

**PERSPECTIVES ON WHAT TO BELIEVE
THE INFORMATION-SENSITIVITY OF THE DOXASTIC 'SHOULD' AND
ITS IMPLICATIONS FOR NORMATIVE EPISTEMOLOGY**

Sebastian Josef Albrecht Becker

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



2016

**Full metadata for this item is available in
St Andrews Research Repository
at:**

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

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

<http://hdl.handle.net/10023/9890>

This item is protected by original copyright

**This item is licensed under a
Creative Commons Licence**

Perspectives on What To Believe

The Information-Sensitivity of the Doxastic ‘Should’
and Its Implications for Normative Epistemology

Sebastian Josef Albrecht Becker



University of
St Andrews

This thesis is submitted in partial fulfilment for the degree of
Doctor of Philosophy
at the University of St Andrews

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

19th of October 2016

1. Candidate's declarations:

I, Sebastian Josef Albrecht Becker, hereby certify that this thesis, which is approximately 75,000 words in length, has been written by me, and that it is the record of work carried out by me, or principally by myself in collaboration with others as acknowledged, and that it has not been submitted in any previous application for a higher degree.

I was admitted as a research student in September 2012 and as a candidate for the degree of Doctor of Philosophy in September 2012; the higher study for which this is a record was carried out in the University of St Andrews between 2012 and 2016.

Date 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 Signature of supervisor

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

In submitting this thesis to the University of St Andrews I understand that I am 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. I 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. I 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 publication of this thesis:

PRINTED COPY

a) No embargo on print copy

ELECTRONIC COPY

a) No embargo on electronic copy

Date Signature of candidate Signature of supervisor

Please note initial embargos can be requested for a maximum of five years. An embargo on a thesis submitted to the Faculty of Science or Medicine is rarely granted for more than two years in the first instance, without good justification. The Library will not lift an embargo before confirming with the student and supervisor that they do not intend to request a continuation. In the absence of an agreed response from both student and supervisor, the Head of School will be consulted. Please note that the total period of an embargo, including any continuation, is not expected to exceed ten years. Where part of a thesis is to be embargoed, please specify the part and the reason.

Abstract

This thesis explores the extent to which the doxastic ‘should’ is information-sensitive and the implications of this for a number of debates in normative epistemology. The doxastic ‘should’ is a special case of the deontic modal ‘should’ and occurs in sentences such as ‘You shouldn’t believe everything you read online’. In the recent semantics literature, it has been suggested that the deontic ‘should’ is information-sensitive, meaning that sentences of the form ‘*S* should do *A*’ are relativized to information-states. After a short introductory chapter, I survey the relevant semantics literature in chapter 2 and provide a simplified contextualist semantics for the doxastic ‘should’, according to which the truth-conditions of sentences containing the doxastic ‘should’ vary with the information-state provided by their context of utterance.

In chapters 3 to 6, I discuss the different kinds of information-states the doxastic ‘should’ can be relativized to and how the respective relativization matters for normative epistemology. Chapter 3 argues that the doxastic ‘should’ has a subjective and an objective sense and that this distinction solves the apparent conflict between subjective epistemic norms and the truth norm for belief. Chapter 4 addresses the question of how one should react to misleading higher-order evidence. I propose that two seemingly opposing views on this issue, Steadfastness and Conciliationism, are both correct. In a sense of ‘should’ that is relativized to one’s first-order evidence, one should remain steadfast in the face of misleading higher-order evidence, but in another sense, which is relativized to one’s higher-order evidence, one shouldn’t. In chapters 5 and 6, I argue that when we advise others on what they should believe, we talk about what they should believe in light of their and our joint evidence. Chapter 7 concludes this thesis with a defence of contextualist semantics for the doxastic ‘should’ against truth-relativist challenges.

Acknowledgements

I owe my greatest philosophical debt to Professor Jessica Brown, my primary supervisor. Jessica provided me with fantastic advice and support over the last three and a half years. She is excellent at helping with the nitty-gritty of writing a philosophy paper as well as with the question of how to present and frame one's ideas. Meetings with Jessica always provided me with challenging objections to think about, but at the same time encouraged me in my work and gave me the feeling that I was on the right track. I could not have asked for a better supervisor.

I am also thankful to Andy Egan, who was my secondary supervisor from 2012 until 2014, and to Matt McGrah, who took over from Andy in 2014. Both of them did not only meet up with me during summers when they were in St Andrews, but also sent me detailed feedback on drafts of papers when they were at their home universities, despite being very busy with their work there.

I was lucky enough to spend January to April 2015 at the University of Southern California, thanks to a grant from Arché. My main supervisor there was Ralph Wedgwood, who was very generous with his time and gave me invaluable comments on almost all chapters of my thesis. I benefited greatly from his expertise on contextualism, the semantics of modals, and epistemic norms. Steve Finlay was also so kind to give me feedback on one of my papers.

I had the great privilege to write my thesis as a research student at Arché. This allowed me to discuss my ideas with brilliant colleagues, who all left an imprint on how I think about philosophy in general and about the questions of my thesis in particular. The *Evidence, Justification, and Knowledge* Research Seminar as well as the *Propositions, Indexicality, Time and Tense* Research Seminar introduced me to new philosophical ideas. The sessions of the Semantics Reading Group and the Probability Reading Group brought me up to speed with formal methods which I was not familiar with before. Finally, the Work In Progress Group provided me with incredibly detailed comments on a number of my papers. I am indebted to the following Arché members and visitors for helpful discussions and/or written comments on my papers (apologies for any omissions): Marvin Backes, Derek Ball, Mark Bowker, Matthew Cameron, Herman Cappelen, Laura Celani, Stew Cohen, Aaron Cotnoir, Alexander Dinges, Daniel Drucker, Claire Field, Noah

Friedman-Biglin, Patrick Greenough, Josh Habgood-Coote, Katherine Hawley, Nick Hughes, Patrik Hummel, Torfinn Huvenes, Ryo Ito, Bruno Jacinto, Poppy Mankowitz, Matt McKeever, Andrew Peet, Patrick Shireff, Justin Snedegar, Wes Skolits, Fenner Tannswell, Ravi Thakral, Caroline Touborg, Brian Weatherson, Alexander Yates, and Julia Zakkou. Furthermore, many thanks to Lynn Hynd for her always reliable and helpful administrative support.

Outside of Arché, I was fortunate to be able to present my work in the St Andrews Philosophy Graduate Research Seminar and at conferences in Aberdeen, Edinburgh, Osnabrück, Pittsburgh, and Stockholm. I thank the audiences at these conferences and in particular my commentators, Jimmy Goodrich, Ulf Hlobil, and Martin Smith. I'm furthermore grateful to the following philosophers inside and outside of St Andrews who helped to improve this thesis through insightful discussions and/or written comments on my papers: Marcia Baron, Matthew Chrisman, Philip Ebert, Renaud-Phillipe Garner, John MacFarlane, Brian McElwee, Robin McKenna, Panos Paris, Peter Sullivan, and Robert Trueman. Finally, a self-standing, longer version of section 5.8 has been published under the title 'Resolving Turri's Puzzle about Withholding' in *dialectica* (vol. 70, issue 2). I am thankful to anonymous referees for *dialectica* for their extremely helpful suggestions.

I would like to thank the St Andrews/Stirling Philosophy Graduate Programme for its financial support in form of a Doctoral Scholarship.

Writing a PhD thesis is a challenging experience. I could always count on my friends to support me in times when things weren't going exactly as they were supposed to, but could also celebrate the successes with them. Besides those friends who I have already mentioned, I want to thank here in particular Tim Dorlach and my brother, Johannes Becker. I owe my greatest emotional debt to my partner, Skaiste Anuzyte. She was with me, either in person or via Skype, every day of the last two years of this project. Her enduring and unconditional support, understanding, and warmth made these two years much lighter and happier than they would have otherwise been.

Finally, I want to thank my parents, Monika and Josef Becker. I cannot express my gratitude for their unwavering financial and emotional support. By creating a culture of open debate at home, they instilled in me the intellectual curiosity that drew me into philosophy. I dedicate this thesis to them.

Contents

1	Introduction	1
1.1	The Setting	1
1.2	Overview of Chapters	4
1.3	Background Assumptions	7
2	The Semantics of the Doxastic ‘Should’	13
2.1	Introduction	13
2.2	Modality in Language	15
2.3	Kratzer’s Standard Semantics for ‘Must’	16
2.4	Semantics for the Deontic ‘Should’	21
2.5	Information-Sensitivity	25
2.5.1	The Phenomenon	25
2.5.2	Option 1: Information-Sensitive Ordering Sources	29
2.5.3	Option 2: DP Models	33
2.6	Semantics for the Doxastic ‘Should’	37
2.6.1	Doxastic Ordering Sources	37
2.6.2	Information-Sensitivity of the Doxastic ‘Should’	40
2.7	A Simplified Semantic Account of the Doxastic ‘Should’	45
2.8	Metasemantics	50
2.9	Conclusion	55
3	The Subjective and the Objective Doxastic ‘Should’	57
3.1	Introduction	57
3.2	The Subjective/Objective Distinction	58
3.3	The Objective Doxastic ‘Should’	62
3.4	Subjective Epistemic Norms and the Truth Norm	68
3.4.1	Subjective Epistemic Norms	68

3.4.2	The Truth Norm	75
3.4.3	The Conflict	77
3.4.4	Resolving the Conflict	79
3.5	Four Objections	80
3.6	Conclusion	87
4	Higher-Order Evidence	89
4.1	Introduction	89
4.2	What Is Higher-Order Evidence?	91
4.3	Motivating Steadfastness and Conciliationism	93
4.4	The Evidence-Ambiguity Theory	97
4.4.1	The Semantics of the Doxastic ‘Should’	97
4.4.2	Approach 1	98
4.4.3	Approach 2	100
4.5	So What Should I Believe?	105
4.6	A Contextualist Model for Theoretical Deliberation	107
4.7	Dilemma or Ambiguity?	113
4.8	Conclusion	116
5	The Collective ‘Should’ and Epistemic Advice	117
5.1	Introduction	117
5.2	What is Epistemic Advice?	118
5.3	Objectivism	125
5.4	Subjectivism	130
5.5	Available Evidence	134
5.6	Advisor’s Evidence	136
5.7	Collectivism	140
5.8	Resolving Turri’s Puzzle about Withholding	142
5.9	Conclusion	149
6	Precisifying Collectivism	151
6.1	Introduction	151
6.2	How to Pool Evidence	152
6.3	Group Justification and Group Knowledge	156
6.4	Is Epistemic Advice Factive?	161
6.4.1	The Factivity Necessity Condition	161
6.4.2	The Factivity Sufficiency Condition	163

6.4.3	Two Arguments against Factivity Intuitions	166
6.4.4	Explaining Away Factivity Intuitions	167
6.5	Conclusion	170
7	The Relativist Challenge	173
7.1	Introduction	173
7.2	Relativism	175
7.3	The Doxastic Integration Problem	178
7.3.1	The Practical Integration Problem	179
7.3.2	Extending the Practical Integration Problem	182
7.3.3	Solving the Doxastic Integration Problem	185
7.4	The Semantic Assessment Problem	189
7.4.1	Disagreement	189
7.4.2	Reports	193
7.4.3	Retraction	197
7.5	Conclusion	200
8	Conclusion	201

Chapter 1

Introduction

1.1 The Setting

Analytic epistemology has for a long time been concerned with knowledge and the semantics of ‘knows’. As looking at the syllabus for almost any introductory epistemology course shows, the question ‘What is knowledge?’ has been central to analytic epistemology in the second half of the 20th century and the beginning of the 21st. There is, however, another way to look at epistemology, according to which the question ‘What should I believe?’ is at the heart of it (Berker 2013). Construed in this way, epistemology is a deeply normative subject. Of course, the question of whether we know possibly has wide-ranging normative implications, as the literature on the norms for practical reasoning or assertion has shown.¹ Furthermore, the concept of knowledge itself is often described as evaluative, similar to thick value concepts like courage or honesty (Ridge 2007: 98). The question of what we should believe is more straightforwardly normative, though, and makes epistemology akin to ethics. Ethics, roughly, is concerned with which action one should choose, whereas epistemology, on this picture, is concerned with which doxastic attitudes one should adopt.

The question of which doxastic attitudes one should adopt is, without a doubt, very important. Our actions matter, but so do our thoughts. Which beliefs we hold, how we represent the world, influences how we navigate in it. However, we do not only care about our beliefs because of their impact on how we act.

¹See, among many others, Williamson (2000), Fantl and McGrath (2002), Hawthorne (2004), and Stanley (2005).

True, justified, or rational beliefs also seem to have intrinsic value for us. Thus, independent of whether the question of what one should believe is at the heart of epistemology or not, it certainly merits our attention.

Even though the question has not been at the centre of epistemology, it has not been ignored either. A famous treatment of it is Clifford's (1877/1999) programatically named 'Ethics of Belief', in which he famously states that "it is wrong always, everywhere, and for anyone to believe anything on insufficient evidence" (ibid.: 77). Bayesian epistemology can be aptly described as being concerned with determining which credence one should adopt in light of one's evidence (Weisberg 2011: sec. 2.5). Furthermore, on a deontic reading of 'justification'², the debate between internalists and externalists about justification is about what one is permitted to believe. Finally, recent discussions about what the norm for belief is, e.g., justification, knowledge, or truth, have moved the question of what we should believe more into the spotlight of current epistemological theorizing.³

These examples from the literature should make clear that I am concerned with the question of what one should believe according to epistemic norms, not according to practical norms. If I offer you one million pounds to believe that there is an elephant in the room, you probably have practical reasons to believe this (Foley 1987: 211). However, you wouldn't have *epistemic* reasons to believe this, presumably because you have clear evidence that there is no elephant in the room. It is notoriously difficult to define what exactly 'epistemic' means (Cohen ms). For our purposes, it should suffice to say that epistemic norms are concerned with which doxastic attitude one should believe in light of considerations pertaining to truth, knowledge, justification, evidence, cognitive reliability, and the like. Practical norms, on the other hand, are about what one should do in light of practical considerations such as money, happiness, or what is morally right. They can be divided into prudential norms, which are about about one should do to further one's interest, and moral norms. Let's call 'should' where it expresses epistemic norms *doxastic* and where it expresses practical norms 'practical'.⁴

²This reading is controversial, though—see Alston (1988) and section 1.3 further below.

³For more on this debate, see chapter 3.

⁴It might be a bit confusing that I use the label *doxastic* to refer to 'should' where it expresses epistemic normativity, and don't speak of the *epistemic* 'should' instead. I do so to avoid the impression that I'm talking about the epistemic modal 'should', as in 'They should be here by now'. Chrisman (2008) uses the same terminology.

Some argue that we do not only have epistemic obligations to adopt certain doxastic attitudes, but that we can also have epistemic obligations to perform certain actions, like gathering evidence.⁵ While I do not want to deny that there are such norms, I focus in this thesis entirely on epistemic norms pertaining to doxastic attitudes, both the coarse-grained ones of belief, disbelief and suspension of judgement, as well as credences.

Given the importance of the question of which doxastic attitude one should adopt for epistemic reasons, it is surprising that there is not much work on the semantics of sentences of the form ‘Subject *S* should have doxastic attitude *D* to proposition *P*’, where the ‘should’ expresses epistemic norms. Since one of the key methods of analytic philosophy is to analyse the meaning of philosophically interesting expressions, looking into the semantics for the doxastic ‘should’ is something normative epistemologists should consider worthwhile. Lack of such research is even more surprising given that the semantic analysis of ‘*S* knows *P*’ has played such a big role in epistemology in the last 20 or so years. Philosophers have defended invariantist⁶, contextualist⁷, contrastivist⁸, relativist,⁹ and expressivist semantics¹⁰ for ‘knows’. In ethics, research into the meaning of ‘should’ or ‘must’ has also been put to good use.¹¹

The present thesis is a contribution to filling this gap in the literature. I believe that a thorough analysis of the semantics of the doxastic ‘should’ can gain interesting insights into a number of debates in epistemology. For example, within Kratzer’s (1991) standard semantics for deontic modals like ‘should’ or ‘must’, it is relatively easy to model a distinction between the doxastic ‘should’ and the practical ‘should’. This is relevant for the question of what the relation between epistemic and practical normativity is.¹² Furthermore, the work that has been done on the semantic distinction between ‘should’ and ‘must’ could throw some

⁵See, for example, Kornblith (1983), Booth (2006, 2009), Pollock (2008: 264), Field (2009: 259), and Goldman (2010: 18).

⁶See, for example, Unger (1984/2002) and Brown (2010)

⁷See, for example, Cohen (1986) and DeRose (1992).

⁸See, for example, Schaffer (2004).

⁹See MacFarlane (2005).

¹⁰See, for example, Chrisman (2006) and Ridge (2007).

¹¹See, for example, Sloman (1970), Jackson (1991), Kolodny and MacFarlane (2010), Björnsson and Finlay (2010), and Finlay (2014).

¹²For discussion of this question see, among others, Foley (1987: 210f.), Raz (2009), Littlejohn (2012), Whiting (2014), and Worsnip (2015a).

light on whether there is something like epistemic supererogation, i.e., something that one should, but not must, believe.¹³

However, I will focus on another semantic feature of ‘should’: its information-sensitivity. This feature has received a lot of attention in the literature on the semantics of deontic modals, in particular in discussion of the so-called miners paradox.¹⁴ That ‘should’ is information-sensitive means that the semantic evaluation of sentences containing ‘should’ can be dependent on contextually provided information-states. According to contextualist semantics, this means that the truth-conditions of a sentence of the form ‘*S* should do action *A*’ vary with the information-state provided by the context in which it is uttered. As a result, such a sentence can, for example, in one context mean that *S* should do *A in light of S’s evidence*, and in another that *S* should do *A in light of the facts*. The most prominent instance of this information-sensitivity of ‘should’ is the distinction between the subjective and the objective ‘should’, which is well-known in ethics. However, some have argued that ‘should’ can be relativized to information-states other than just the subject’s evidence and the facts—for example to the combined evidence of speakers and their addressees.

The core idea of this thesis is to investigate to what degree the doxastic ‘should’ is information-sensitive and to explore the implications of this for a number of debates in epistemology.

1.2 Overview of Chapters

The discussion of information-sensitivity in the semantics literature focuses almost exclusively on cases where ‘should’ expresses practical norms.¹⁵ The *practical* ‘should’ and the doxastic ‘should’ are both instances of the deontic ‘should’. In chapter 2, I explore new territory when I examine to what degree the semantic models of the information-sensitivity of the deontic ‘should’ presented in the literature, which have been developed with the practical ‘should’ in mind, also work for the doxastic ‘should’. Looking at the proposed contextualist models

¹³Hedberg (2014) and McElwee (ms) discuss, and positively answer, this question.

¹⁴See Kolodny and MacFarlane (2010), Björnsson and Finlay (2010), Dowell (2013), Silk (2014), and Willer (forthcoming).

¹⁵The only exceptions are MacFarlane (2014: 298) and Wedgwood (2007: ch. 5.2), who mention this *en passant*.

on the market, I will argue that Silk (2012), Dowell (2013), and Carr's (2015) models can account for the information-sensitivity of the doxastic 'should', but Cariani et al. (2013) and Charlow's (2013) models cannot. I close by presenting a simplified version of Silk (2012), Dowell (2013), and Carr's (2015) semantics, according to which the truth-conditions of sentences of the form ' S should have D to P ' vary with contextually provided information-states and, roughly, mean ' S should have D to P in light of i ', where i is the information-state provided by the context of utterance. I call this position *Doxastic Contextualism*.

The idea that the doxastic 'should' is information-sensitive is almost completely novel in epistemology, and one of my main jobs in this thesis is hence to provide evidence for Doxastic Contextualism. I will do so by pointing towards certain practices of epistemic evaluation that I think are best explained by the information-sensitivity of the doxastic 'should', relying on our linguistic intuitions and on metasemantic considerations. However, that the doxastic 'should' can be relativized to different information-states is in and of itself not exciting news for the epistemologist. The reason it matters for epistemology is that it can solve a number of puzzles and apparent conflicts in recent debates, or so I will argue.

In chapter 3, I address the only instance of the information-sensitivity of the doxastic 'should' that has received attention in epistemology: the subjective/objective distinction. The idea that what one should believe is somehow sensitive to one's body of evidence is almost uncontroversial among epistemologists. I will present arguments that we also have a practice of talking about what one should believe in the objective sense of 'should', i.e., in light of the facts. The implication of this, *contra* Gibbons (2013) and others, is that the conflict between the truth norm for belief and more subjective norms, like the justification norm for belief, is only apparent since the first employs the objective sense of the doxastic 'should' in its formulation and the second the subjective sense.

In chapter 4, I will address a further kind of information-sensitivity, the distinction between what one should believe in light of one's first-order evidence and what one should believe in light of one's higher-order evidence. This has interesting implications for the peer disagreement debate, as well as the more general question of how one should react towards misleading higher-order evidence. My claim is that those who argue that one should remain steadfast in light of peer disagreement or other kinds of misleading higher-order evidence have this

first sense of the doxastic ‘should’ in mind (the one relativised to one’s first-order evidence), whereas those who argue that one should be conciliatory and adjust one’s doxastic attitude in light of misleading higher-order evidence employ the second sense. The upshot is that we neither have to choose between steadfastness and conciliationism in the peer disagreement debate nor between analogous views in the higher-order evidence debate in more general.

In chapter 5, I turn to what I call the collective sense of the doxastic ‘should’. I argue that in some contexts we talk about what a subject should believe in light of the collective evidence, i.e., in light of their, i.e., the subject’s,¹⁶ and the speaker’s combined evidence. We do so when we are giving epistemic advice, which is, roughly, the practice of helping subjects who are deliberating about what to believe. I provide an account of this practice, which has so far not been discussed in the literature. Furthermore, I propose and argue for **Collectivism**, the view that the collective doxastic ‘should’ is the one employed for this purpose, rather than other possible senses like the subjective or objective doxastic ‘should’. This tells us under which conditions epistemic advice is correct and thus helps us to gain a better understanding of this hitherto under-researched practice. A more direct implication of the insight that the doxastic ‘should’ can be relativised to the collective evidence is that it solves a puzzle described by Turri (2012) in his paper ‘A Puzzle about Withholding’, which he declares to be unsolvable. Chapter 6 discusses **Collectivism** further and explores what exactly collective evidence is; whether epistemic advice of the form ‘You should believe *P*’ is true only if *P* is true; and examines possible connections between **Collectivism** and the literature on group justification.

I hope that I will have provided convincing evidence for Doxastic Contextualism in chapters 3 to 5. However, contextualist semantics for a number of terms, like epistemic modals, ‘knows’, taste predicates, and deontic modals, have faced serious criticism, especially from defenders of relativist semantics, but also from those arguing in favour of invariantist semantics for these terms. Since chapters 3 to 5 should have severely undermined invariantism about the doxastic ‘should’, I will, in chapter 7, defend Doxastic Contextualism against those challenges against the backdrop of relativism. The first objection is that contextualist semantics of the deontic ‘should’ can’t account for the practice of advice. I

¹⁶Here and throughout this thesis I use the singular ‘they’ to refer to indeterminate persons.

argue that Dowell’s (2013) response to this objection, put forward in defense of a contextualist semantics for the practical ‘should’, does not work for the doxastic ‘should’, and I suggest an alternative reply to defend Doxastic Contextualism. The second objection is that contextualists struggle to explain linguistic phenomena like disagreement, attitude reports, and retractions. I will offer different responses, arguing for some of these phenomena that the data relativists provide is not as convincing as it seems at first sight and for other phenomena that they are caused by the semantic blindness of ordinary speakers.

1.3 Background Assumptions

Before we begin, I need to highlight two major assumptions I make in this thesis. The first is that there are epistemic norms or, at the very least, that sentences of the form ‘*S* should (doxastically) believe *P*’ can be true. The second assumption is that I work with a truth-conditional semantics for the doxastic ‘should’. In the following, I will motivate these assumptions.

As I explained, I take the doxastic ‘should’ to express epistemic norms. If there are no such norms, one could argue that there is also no such ‘should’ and that this thesis is accordingly pointless. Now, there seem to be three kinds of reasons why people might reject the idea that there are epistemic norms. First, Papineau (2013) argues in his aptly named ‘There Are No Norms for Belief’ that there is no *sui generis* epistemic normativity. He admits that it is true that when we have the aim of having a true belief about a proposition *P*, we should do things that help us achieve this aim, e.g., believe what is supported by our evidence. However, when we don’t pursue this truth goal with respect to a proposition *P*, then our doxastic attitude to *P* is not subject to any epistemic normative constraints. He gives the example of John, who has untreatable cancer and evidence that this is the case, but who maintains the false belief that he is fine and thereby avoids the emotional pain that would come with having the belief that he’s terminally ill. Papineau (2013: 68) thinks that John “has done nothing wrong at all” since holding a true belief about his health is not in his interest. He claims that this shows that there is no real epistemic normativity and that what we should believe is rather determined by prudential considerations, not epistemic ones. If there was epistemic normativity, it would be categorical and demand to be followed no

matter what our interests are.¹⁷

One way to reply to this is to contest Papineau's (2013) judgement and insist that John should not believe that he's fine, even if it makes him happy. Kelly (2003) considers cases similar to the one of John, where people do not care about having true beliefs, as counterexamples to epistemic instrumentalism. Epistemic instrumentalism claims that epistemic reasons derive their normative force from epistemic goals, like the goal to have true beliefs. Kelly (2003) argues that epistemic instrumentalism is mistaken precisely because we have epistemic reason to believe what our evidence supports even if we don't care to have a true belief about the proposition in question. Thus, Kelly (2003) would argue that John should believe that he has cancer and that this shows that epistemic instrumentalism is wrong, and not that there are no epistemic norms.

Papineau (2013: 69) admits that there is a practice of epistemically evaluating others, but contests that this implies that there is true epistemic normativity. Given this, he could reply to Kelly (2003) that indeed, we might *say* of John that he should believe that he has cancer, but that this doesn't mean that John should *really* believe this. What he really should believe is that he is fine since this would make him happy. I think there are two possible replies to this. First, one could argue that in a case like John's there is no such thing as what he *really* should believe. Rather, prudentially, in terms of what best serves his interest, John should believe that he's healthy, and doxastically, John should believe that he's ill. On this view, there are prudential and epistemic norms and one kind doesn't trump the other. Thus, something can be a proper norm or be 'genuinely' normative without being such that it categorically prevails over other considerations. Some argue that prudential reasons can stand in a similar relation to moral reasons. It can be in my interest to do something that is morally wrong. According to a position often attributed to Sidgwick, Henry (1874), in such situations there is nothing we really should do or all-things-considered. Prudentially, we should do

¹⁷In his (2010), Nelson argues that there are no positive epistemic duties, i.e., that we can be permitted to hold a belief, but never obliged to have one. If this was true, it would be a threat to my project since I assume that sentences of the form 'You should believe *P*' can be true. He rests his argument on the assumption that when we positively should believe something this is only so because we have a prudential interest in having a true belief about the subject matter at hand (ibid.: 87). He concludes that when we should believe something this is hence not really due to purely epistemic reasons. Accordingly, his position is basically the same as Papineau's (2013) and what I write in the following applies to him as well.

what's in our interest, but morally we should do what morality requires, and no more can be said on this.

According to a second, more concessive reply, it is true that John *really* or *all-things-considered* should believe that he's healthy. However, doxastically he should believe that he has cancer, but it just happens that prudential norms trump epistemic norms. On this view, epistemic norms are maybe more akin to norms of etiquette and a tier below 'real' norms like moral or prudential norms. Papineau (2013) would take this to be basically his point: epistemic norms aren't truly normative. While I don't favour this more concessive position (my sympathies rather lie with the first reply), it wouldn't make this thesis a useless investigation if it turned out true. According to this position, sentences of the form '*S* should have *D* to *P*' can be true, even though they don't express 'proper' norms, just like the sentence 'You should eat with fork and knife'. My investigation of the truth-conditions of sentences of the form '*S* should have *D* to *P*' would still help to resolve the kind of puzzles epistemologists have got themselves into by overlooking the fact that the doxastic 'should' is information-sensitive, even though some would consider these debates less interesting than expected since they are only about second-tier norms.

A second argument against epistemic norms refers to the principle that 'ought' (or 'should') implies 'can'. Alston (1988) famously argued that since we don't have control over our beliefs we can't apply deontic concepts like 'obligation' or 'permission' to our beliefs. If I cannot help but believe *P*, it can't be right that I should believe $\neg P$ (or *P*).

This argument can be rebutted in at least two ways. First of all, some defend direct doxastic voluntarism, i.e., the view that we have direct control over our beliefs. For example, Ginet (2001) maintains that we can decide what to believe in situations where our evidence is inconclusive. Montmarquet (1986) and Steup (2000) claim that where our belief-formation is purely led by our evidence, our belief should be considered as having been formed freely, just as actions that are guided by good practical reasons should be considered free actions. Second, others have denied that the doxastic 'should' implies 'can'. Feldman (2000) puts forward the idea that the doxastic 'ought' (or 'should') is a so-called role-'ought'. For example, even if a bad biology teacher is not capable of explaining natural selection clearly to his students, he should do so nonetheless, qua his role as

biology teacher. Similarly, if due to wishful thinking a mother cannot but believe that her son did not commit a murder, even though her evidence clearly supports this, she should nonetheless, qua her role as a believer, believe that he did commit it. Chrisman (2008), in a somewhat similar spirit, suggests that the doxastic ‘ought’ (or ‘should’) is an ‘ought-to-be’, which describes what would be best, rather than an ‘ought-to-do’, which ascribes obligations to agents, and that only this latter ‘ought’ implies ‘can’. Given that there is apparently quite a bit of reasonable resistance to the idea that there are no epistemic norms because we can’t control our beliefs, this argument is therefore hardly sufficient to discredit the project of this thesis.

Finally, metaethical anti-realists will deny that there are epistemic norms because they deny that there are any norms at all. For example, Olson (2011) points out that Mackie’s (1977) famous argument from queerness equally applies to epistemic norms. I don’t think though that I am necessarily committed to some form of metaethical (or better: metanormative) realism. Expressivists like Blackburn (1993) and Gibbard (2003) have propagated ‘quasi-realism’, the view that strictly speaking there are no normative properties, but that normative sentences can nonetheless be true and false, for example by advocating a deflationary account of the truth-predicate. As I’ve pointed out above, this dissertation can make a valuable contribution to the literature even if one assumes that there are no ‘proper’ epistemic norms because they are more like etiquette norms as long as sentences of the form ‘*S* should have *D* to *P*’ can be true or false. The same holds if there are no epistemic norms at all.

The second major assumption I make is that I work with a truth-conditional semantics for the doxastic ‘should’. The main reason for this is convention. As I will explain in a bit more detail in chapter 2, the standard semantics for deontic modals is Kratzer’s, which happens to be developed within a truth-conditional possible world semantics. Given the number of issues in epistemology that I cover in this dissertation, I needed to decide on one semantic framework. It would simply not be manageable to look at the issues I discuss through the lense of all possible semantics accounts of the doxastic ‘should’.

There are good reasons why Kratzer’s is the standard account. The possible world semantics it is formulated in has become standard in formal semantics since it can model the compositionality of meaning particularly well, i.e., how complex

linguistic expressions derive their meaning from the meaning of their parts (Heim and Kratzer 1998). Not all alternative semantic accounts of deontic modals can claim this. Expressivism, which I just mentioned, has traditionally been interpreted as a semantic theory, according to which the meaning of normative sentences consists in the non-cognitive attitudes they express. The Frege-Geach problem is just an instance of expressivism's more general difficulty with accounting for compositionality, and hence this standard version of expressivism has become less popular.¹⁸ Recently, some authors have therefore come to propose views that are expressivist in spirit, but aren't semantic accounts of normative language. An example of this is Ridge (2014). According to his 'ecumenical expressivism', sentences of the form '*S* should do *A*' mean 'In light of all acceptable standards, *S* should do *A*' and one's normative attitudes determine what is meant by 'acceptable'. On this view, the meaning of normative sentences is not the non-cognitive attitudes they express, but these attitudes rather fix the truth-conditional content of such sentences. This makes ecumenical expressivism a metasemantic, not a semantic theory. Ridge (2014) relies on Kratzer's account himself, so ecumenical expressivism is perfectly compatible with assuming a truth-conditional semantics for the doxastic 'should'.

However, there are indubitably semantic theories for deontic modals that aren't truth-conditional but can handle compositionality. A non-truth-conditional theory that is finding more and more recognition among philosophers is dynamic semantics. It captures the meaning of sentences not in terms of their truth-conditions but in how they affect the discourse if they are accepted by the discourse participants (Yalcin 2012). Originally developed to deal with anaphora (Kamp 1981; Heim 1982), it has very recently been applied to deontic modals, too. Marra (2014) proposes that the meaning of sentences containing deontic modals consists in the potential change they have on 'ideal world relations' in the context of the discourse. These relations are roughly the normative judgements accepted by the discourse participants and the meaning of should-sentences is accordingly the change they bring to these judgements if they are accepted. Starr (forthcoming) combines expressivism and dynamic semantics and suggests that the meaning of normative sentences is the potential effect they have on the (non-cognitive) preferences of the discourse participants.

¹⁸See Schroeder (2008) for a comprehensive criticism of expressivism as a semantic theory.

I can't offer here a satisfying comparison of the advantages and disadvantages of dynamic semantics and the truth-conditional semantics I will rely on throughout this dissertation. As explained above, I need to focus on one semantic account of the doxastic 'should' if I want to cover all the debates in normative epistemology that I think can benefit from checking what the information-sensitivity of the doxastic 'should' implies for them. So rather than looking at one issue in normative epistemology with the help of a number of semantic theories, I have decided to look at a number of issues in epistemology with the help of one semantic theory. Effectively, this is a dissertation in epistemology that utilizes semantic research rather than a dissertation in semantics. Given that I thus had to choose one semantic theory, it seems right to go with the standard one.

That said, Willer (forthcoming: sec. 3.2) has argued that dynamic semantics can account well for the information-sensitivity of deontic modals. Thus, the phenomenon I am analysing here is not one that only truth-conditional semanticists recognize. Which implications for normative epistemology could be drawn from a dynamic semantic analysis of the information-sensitivity of the doxastic 'should' is a question I can't answer here and leave for further research.

To conclude, my assumption that there are epistemic norms, or at least that sentences of the form '*S* should have *D* to *P*' can be true, is reasonable, and so is my assumption that the semantics for the doxastic 'should' is truth-conditional.

Chapter 2

The Semantics of the Doxastic 'Should'

2.1 Introduction

This thesis investigates the implications of the information-sensitivity of the doxastic 'should' for current debates in normative epistemology. The doxastic 'should' occurs in sentences of the form 'Subject S should adopt doxastic attitude D to proposition P '. This 'should' is a kind of the more general deontic 'should', which is a normative modal, roughly expressing what would be best or how things ought to be. How things should be can be evaluated along different dimensions. For example, the practical 'should' concerns what we should do in light of practical norms. The doxastic 'should' expresses a different normative dimension, that of epistemic normativity. What the practical and the doxastic 'should' have in common is that they are instances of the deontic 'should'.

The information-sensitivity of the doxastic 'should' is the phenomenon that a sentence like ' S should (doxastically) have D to P ' is implicitly relativized to information-states. On the contextualist semantics that I will put forward in this chapter, this phenomenon is taken to be that such a sentence means, roughly, 'in light of i , S should have D to P ', where i is the information-state provided by the context in which the sentence is uttered. One could accept this semantics, but posit that contexts of utterance always provide the same kind of information-state, e.g., the subject's. Contrary to this, I argue in this thesis that the doxastic 'should'

is fully information-sensitive and that contexts of utterance provide different kinds of information-states. I call this view *Doxastic Contextualism*.

This chapter does two things. Several authors have tried to give a semantics for the deontic ‘should’ that captures its information-sensitivity. However, all of the suggested models focus on occurrences of the deontic ‘should’ where it carries a practical sense. The first important contribution of this chapter is that I show that Dowell (2013), Silk (2014) and Carr’s (2015) semantics can also account for the information-sensitivity of the *doxastic* ‘should’. It is thus not against the ‘general wisdom’ of current semantic theory to posit that the doxastic ‘should’ is context-sensitive to information-states. However, I also show that Cariani et al. (2013) and Charlow’s (2013) alternative models face some difficulties with this. An interesting upshot of my discussion is consequently that it gives semanticists one more parameter by which to judge semantic models of information-sensitivity: as I will provide evidence in the next chapters that the doxastic ‘should’ is information-sensitive, the failure of Cariani et al. (2013) and Charlow’s (2013) semantics to capture this speaks against them and in favour of the alternatives.

The second important thing I do in this chapter is to present a simplistic contextualist semantics for the doxastic ‘should’ that accounts for the latter’s information-sensitivity. My account is inspired by and in line with Dowell (2013), Silk (2014) and Carr’s (2015) models, but lacks their detail. According to it, a sentence of the form ‘*S* should have *D* to *P*’ is true iff having *D* to *P* is supported by the information-state *i* provided by the context in which the sentence is uttered.

To be clear, this chapter is not supposed to offer an argument in favour of Doxastic Contextualism. Rather, the chapter explains Doxastic Contextualism and shows how it can be integrated into the current literature on the semantics of deontic modals. I will provide evidence for Doxastic Contextualism in chapters 3 to 5. However, as I have stated above, a contextualist semantics is only one possible way to take account of the information-sensitivity of the doxastic ‘should’. As I suggested towards the end of chapter 1, dynamic semantics might also have a way to account for the phenomenon that the doxastic ‘should’ is sensitive to information-states and the evidence provided in chapters 3 to 5 could thus equally be evidence for a dynamic semantics of the information-sensitivity of the doxastic ‘should’. For reasons of scope, I won’t assess such an account, or any other

non-truth-conditional one, here and don't claim that Doxastic Contextualism is superior to it. That said, I will discuss one rival account of the information-sensitivity of the doxastic 'should', namely relativism, in chapter 7.

I will proceed as follows in the present chapter. Section 2.2 gives a quick overview of modal expressions, the semantic category that 'should' falls into. Section 2.3 explains Angelika Kratzer's standard semantics for modal expressions, in particular for the term 'must'. Section 2.4 discusses several recent attempts at developing Kratzer's account of 'must' into a semantics for the deontic 'should'. Section 2.5 looks at the mentioned current contextualist semantic accounts of the information-sensitivity of the deontic 'should', and in section 2.6 I show that some, but not all of them, can capture the information-sensitivity of the doxastic 'should'. Section 2.7 presents my simplistic semantics for the information-sensitivity of the doxastic 'should'. In section 2.8, I turn to metasemantics and discuss which features of a context determine what the information-state provided by this context is.

2.2 Modality in Language

'Should' is a modal auxiliary like 'must' or 'can'. It belongs to the big family of modal expressions. Besides modal auxiliaries, there are modal verbs like 'have to' or 'need to', modal adjectives like 'possible' or 'necessary', and modal nouns like 'possibility' or 'necessity'. There exist further ways to express modality in English, like modal adverbs, verbal mood, and many others (Portner 2009: 4-8). What all these expressions have in common is that they allow us to talk about "situations which need not be real" (ibid.: 1). For example,

- (1) It is possible that the Queen will become a 100 years old.

is true, even though it might turn out wrong that the Queen will live to her 100th birthday. This shared characteristic of modal expressions is also called "displacement" (von Stechow and Heim 2011: 2).

A central feature of modality is that there are different "flavours" of it (Kratzer 1977). This is nicely illustrated by the following two sentences:

- (2a) They must be in London by now as they left Edinburgh quite early in the morning.

(2b) You must honour your parents.

Both sentences contain ‘must’, but ‘must’ does not mean the same in them respectively. The ‘must’ in (2a) is the one of epistemic modality, meaning something like ‘certainly’. The ‘must’ in (2b), however, is the one of deontic modality, roughly meaning ‘morality requires that’. There are further modal flavours. Deontic modality is a subtype of what some call priority modality. All subtypes of priority modality concern what is best or preferred in terms of certain rankings: deontic modality is about what is best in light of rules, like those of morality. Bouletic modality concerns rankings based on one’s desires (2c), and teleological modality what is best given one’s goals (2d).¹

(2c) I must see this!

(2d) If you want to go to Harlem, you have to take the A train.

A third major group of modals besides epistemic and priority modals are circumstantial or dynamic modals, which express what is possible given one’s abilities:

(2e) Mike can’t sing.

To sum up, a striking characteristic of modal expressions like ‘should’ is that they don’t have a fixed meaning, but can express different dimensions of modality. We will now turn to the standard semantics for modal terms by Angelika Kratzer, which provides a formal framework that captures this central feature of modal expressions.

2.3 Kratzer’s Standard Semantics for ‘Must’

In a series of articles (1977, 1981, 1991), Kratzer gives a semantics for modal terms that is often described as the standard one. According to Portner (2009), it deserves this title since it is the “working assumption in the semantics research” and the theory “which a linguistically oriented semanticist is most likely to recommend to students or colleagues who wish to learn about the theory of

¹Different authors have suggested different taxonomies for types of modality. I am here following Portner (2009: 135). I am not committed to his taxonomy. My aim is just to display the diversity of modal flavours.

modality” (ibid.: 47). Bronfman and Dowell (forthcoming) praise it as the “view to beat” due to its “simple and highly unified explanations of a wide range of language use”. Finlay (forthcoming) describes it as being “widely regarded as orthodoxy.”

Kratzer herself does not provide an account for ‘should’ (Finlay 2014: 77, n. 64), but does so for the related ‘must’. In section 2.4, I will discuss semantics specifically provided for ‘should’. These semantics are further developments of Kratzer’s account of ‘must’ by other authors or can at least better be explained in contrast to it. For this reason, I will here give a short explanation of Kratzer’s semantics for ‘must’ first.

Kratzer’s account is a refinement of possible world semantics for modal expressions. According to this semantics, modal expressions are propositional operators. That is, they are functions that take propositions as arguments and assign them truth-values. The proposition they take as argument is called the ‘prejacent’ in the literature. Take the following example:

(3) Peter must pay his taxes.

The logical form of (3) is

(LF3) MUST (Peter pays his taxes),

where the proposition expressed by ‘Peter pays his taxes’ is the prejacent embedded under the modal ‘must’ (von Stechow and Heim (2011: 30)). A central assumption of possible world semantics for modal expressions is that modals are not just any kind of propositional operator, but quantifiers. They take a proposition and map it onto the truth-value 1 depending on whether the proposition holds at worlds in a set of worlds the modal is quantifying over, the so-called ‘modal domain’. Necessity modals like ‘must’ are universal quantifiers and map a proposition P to 1 iff P holds at all worlds in the modal domain. Possibility modals like ‘can’ or ‘might’ are existential quantifiers and map P to 1 iff P holds at some worlds in the modal domain (ibid.). Modals of different flavours quantify over different domains (ibid.: 32). Very roughly, the modal domain of epistemic modals consists of the worlds compatible with the knowledge of the speaker, whereas deontic modals quantify over all the worlds compatible with compliance with the relevant rules

(ibid.: 33f.). We get the following lexical entries for the epistemic and deontic ‘must’:

Deontic Must $\llbracket \text{must}_{deontic} \phi \rrbracket^{c,w} = 1$ iff for all worlds v such that the rules applying to w according to c are complied with at v , $\llbracket \phi \rrbracket^{c,v} = 1$.

Epistemic Must $\llbracket \text{must}_{epistemic} \phi \rrbracket^{c,w} = 1$ iff for all worlds v such that v is compatible with what the speaker of c knows at w , $\llbracket \phi \rrbracket^{c,v} = 1$.

A quick explanation of the notation is in order. The double brackets are interpretation functions that map a linguistic expression onto its semantic value, i.e., in the case of sentences onto truth-values (ibid.: 5). Interpretation functions are relativized to a context and a world. This reflects that sentences have truth-values only relative to a world—a sentence is true at a world—and relative to a context since sentences with indexical terms like ‘I’ or ‘here’ only express a particular proposition relative to a context (ibid.: 7).² **Deontic Must** says that the semantic value of a sentence which consists of a deontic ‘must’ embedding a prejacent ϕ is ‘1’ at a world w iff ϕ holds at all worlds in which the rules applying in w , according to the context c where the sentence is uttered, are complied with. For example, ‘Peter must pay his taxes’ is true in the actual world iff Peter pays his taxes in all worlds where the rules of morality (or the fiscal legal rules), as they apply to the actual world, are complied with.

Kratzer brought two main contributions to this general picture (Portner 2009: 47). First, on the above account, ‘must’ is ambiguous, which means that each ‘must’—the deontic, the epistemic, the bouletic, etc.—has its own lexical entry. Kratzer (1977), on the other hand, provides a unified account, on which ‘must’ has a single lexical entry. Which modal flavour a particular ‘must’ expresses is a function of the context of utterance of the relevant sentence. The set of worlds it quantifies over is provided by what Kratzer calls “conversational backgrounds”, functions given by the context. Which conversational backgrounds are provided by the context determines which sets of worlds ‘must’ quantifies over and thus which modal flavour it expresses. A modal like ‘must’ is therefore on this picture context-

²Moreover, sentences need to be relativized to a third parameter, an assignment function, which assigns values to pronouns, traces, variables, etc. I’ll ignore this here since it is irrelevant for our purposes.

sensitive in a way similar to indexicals. To use Kaplan (1989b)’s terminology, modals have a unified meaning in the sense that they have the same character, a function from contexts to contents, but are context-sensitive since their character assigns different contents to different contexts. The character of ‘I’ maps a context to the speaker of the context. That is, ‘I’ has one character, but its content varies with whoever is the speaker in the given context. Similarly, ‘must’ has one character, which assigns different contexts different functions from propositions to truth-values, depending on which set of worlds is provided by the respective context as the modal domain for ‘must’.

The second contribution is the specific way in which in Kratzer’s model the modal domain is determined. According to Kratzer (1981), context provides two conversational backgrounds, two functions f and g , which take worlds as an argument and assign them sets of propositions. The first function f is called the *modal base*. Depending on the context, it assigns worlds different sets of propositions. An epistemic modal base maps all the propositions known by the speaker in a world onto that world. A circumstantial modal base assigns a world certain facts that hold at that world (Kratzer 1991: 646). Following Carr (2015), I will call the intersection of the propositions assigned by a modal base f , $\cap f(w)$, the *modal background*. The second function g , the *ordering source*, ranks the worlds in the modal background. The ordering source does so by ranking the worlds according to how close they are to the propositions in $g(w)$, i.e., the set of propositions assigned to w by g , which I will call the *ordering background*. For a world v in the modal background to be best according to the ordering background $g(w)$, or short, to be $g(w)$ -best, the following must hold: all propositions in the ordering background that are true at any other world in the modal background are also true at v and there are some propositions in the ordering background that are true at v , but at no other worlds in the modal background. The set of $g(w)$ -best worlds in the modal background is the modal domain.³ Equipped with

³The assumption that every ordering source used by natural language induces a set of best worlds is called the ‘limit assumption’ (Lewis 1973). It is controversial since there are possible ordering sources that do not induce sets of best worlds. For example, think of an ordering source reflecting King Midas’ wish to always be richer. Whatever number of gold coins he owns in a world, there will always be a world in which he owns one more gold coin and which hence will be better (Portner 2009: 65). Kratzer (1981) herself does not make the limit assumption and therefore provides an analysis of ‘must’ which is more complex than the following **Must** in order to accommodate ordering sources that don’t induce sets of best worlds. I ignore this complexity here since it’s irrelevant for our purposes.

this, we get the following lexical entry for ‘must’ (Portner 2009: 67):

Must $\llbracket \text{must } \phi \rrbracket^{c,f,g,w} = 1$ iff for all worlds v such that v is among the $g(w)$ -best worlds in $\cap f(w)$, $\llbracket \text{must } \phi \rrbracket^{c,f,g,v} = 1$.

In other words, ‘must ϕ ’ holds true at a world w iff ‘ ϕ ’ is true at all those worlds in the intersection of the set of propositions assigned by the modal base to w which are $g(w)$ -best. Notice that the interpretation function is now also relativized to a modal base f and an ordering source g . Where these parameters are not determined by some linguistic material—for example by phrases like ‘in light of the British tax laws’—they are determined by context (Bronfman and Dowell forthcoming).

There are several reasons why Kratzer (1981, 1991) sees a necessity for a second conversational background, an ordering source, on top of the modal base to determine the modal domain. At first sight, one might think that a modal base alone would do. For example, this one conversational background could in contexts where the modal has a deontic flavour simply be a function to propositions describing what the rules prescribe, or in context where the modal is epistemic a function to the propositions known by the speaker. However, such a simpler model cannot account for expressions of graded modality such as

- (4) There is a slight possibility that they are not coming tonight.
- (5) Eating vanilla ice cream is more desirable than eating chocolate ice cream.

If we only have a modal base, we can only make coarse-grained distinctions between worlds: between those that are epistemically possible and those that are not or those in which what is desirable is achieved and those in which it is not. But (4) and (5) require finer distinctions. We have to compare worlds in terms of how likely or desirable they are. Since ordering sources can provide rankings along these lines, they help to account for graded modality. A further advantage of two conversational backgrounds is, according to Kratzer (1991: 642), that they can better resolve certain deontic paradoxes involving modals and conditionals, e.g., the Samaritan Paradox, and account for deontic modals expressing inconsistent rules.

2.4 Semantics for the Deontic ‘Should’

‘Should’ and ‘must’ differ in meaning. For one, ‘must’ entails ‘should’, but not vice versa, as the following examples indicate (Harman 1993; von Fintel and Iatridou 2008: 117).

- (6a) You should visit your grandma, but it is not the case that you must visit her.
- (6b) #You must visit your grandma, but it is not the case that you should do this.

Since ‘must’ is thus logically stronger than ‘should’, ‘must’ is often called a *strong necessity modal*, whereas ‘should’ is labelled a *weak necessity modal*. Standardly, philosophers describe the difference in meaning between the deontic ‘must’ and the deontic ‘should’ as follows: ‘You should do *A*’ means that doing *A* is the best of one’s options (where ‘*A*’ is a variable for an action, understood broadly, i.e., it can also stand in for doxastic attitudes). On the other hand, ‘you must do *A*’ means that doing *A* is the only option one has, for example in order to achieve one of one’s goals (Sloman 1970: 390f.; Williams 1981: 125). Due to these differences in meaning, we need a semantics for ‘should’ that is different from that for ‘must’.

The reader might wonder why I focus on the doxastic ‘should’ rather than ‘must’ in this thesis. If the standard semantics for modals, Kratzer’s, gives an account of ‘must’, but not ‘should’, wouldn’t it be easier to investigate what the implications of the information-sensitivity of the doxastic ‘must’ are for normative epistemology? I have three reasons for being concerned with the doxastic ‘should’ instead. First, epistemologists seem to rather talk about what we ‘should’ than what we ‘must’ believe. Since I want this investigation into the semantics of epistemic-normative discourse to yield insights on debates in normative epistemology, it is thus wise to focus on those aspects of discourse which normative epistemologists are more likely to engage in. Second, from a Google search of expressions like ‘You must believe *P*’ it appears that they express more often than expressions like ‘You should believe *P*’ that the addressee should believe *P* for non-epistemic reasons or at least not fully epistemic reasons, but rather for practical reasons, like religious ones. Take the following verses from the New International version of the Bible:

But when you ask, you must believe and not doubt, because the one who doubts is like a wave of the sea, blown and tossed by the wind. That person should not expect to receive anything from the Lord. (James 1:6-7)

An explanation of this could be that since ‘must’ is stronger, it expresses a more stringent form of obligation or normative pressure and that epistemic concerns rather seldom call for such an ‘urgent’ tone, unlike practical concerns. Third, and most importantly, while information-sensitivity in weak deontic necessity modals like ‘should’ has been widely discussed, it seems not to be a feature of strong deontic necessity modals like ‘must’. Since I’m interested in the relevance of this semantic feature for normative epistemology, I focus on ‘should’.

While Kratzer herself does not give a semantics for ‘should’, several authors have done so. Some of these theories are quite straightforward developments of Kratzer’s general framework for modals, where others diverge more radically from it. In the following, I will present these theories, in order to be able to show later in subsection 2.6.1 how the doxastic ‘should’ can be modelled within these current semantic accounts of ‘should’.

First, there are a couple of authors who provide a semantics for ‘should’ that keeps crucial elements of Kratzer’s theory. For example, von Stechow and Iatridou (ms, 2008) suggest to account for the difference between ‘must’ and ‘should’ by simply adding a third conversational background, a second ordering source. Let’s call this secondary ordering source ‘ h ’.⁴ The conversational background h ranks the $g(w)$ -best worlds in $\cap f(w)$. To see how this captures the difference between ‘should’ and ‘must’, take the following examples:

(7a) In order to travel from St Andrews to Edinburgh, you must leave Fife.

(7b) In order to travel from St Andrews to Edinburgh, you should take the train.

⁴These authors, and some of the others whose semantics for ‘should’ I will discuss in the following, actually focus on ‘ought’. However, ‘ought’ and ‘should’ are extremely similar in meaning and what applies to one of them in terms of information-sensitivity will apply to the other, too. Cariani (2013: 73) calls ‘ought’ a “near-synonym” of ‘should’; Finlay (2014: 72) simply describes them as synonymous.

Von Fintel and Iatridou suggest that in the case of such teleological modals the primary ordering source g simply assigns that proposition to a world w that describes the state of affairs mentioned by the ‘in order to’-phrase. In light of this ordering source, those worlds are best in which the subject achieves the goal described by the phrase. Thus, assuming that w is the actual world, it holds that the subject must leave Fife since in all the $g(w)$ -best-worlds in the set of worlds determined by the modal base $f(w)$, the subject leaves Fife.⁵ This captures the sense of ‘must’ as being similar to ‘only’: leaving Fife is the only option for the subject to travel to Edinburgh (in comparison to the option of not leaving Fife). The second ordering source h ranks all these $g(w)$ -best-worlds by the standard of how well the goal is achieved in these worlds, for example, whether it was done in a cheap and comfortable way. If (7b) is true, this means that in all worlds where the subject travels from St Andrews to Edinburgh in the most preferable way, the subject takes the train. This account of ‘should’ aims at modelling the idea that ‘should’ is akin to ‘best’. A further advantage of this semantics is that the set of worlds ‘should’ quantifies over is a proper subset of the set of worlds ‘must’ quantifies over since the former set only contains the $h(w)$ -best worlds in the latter set. It follows that ‘must’ entails ‘should’, but not *vice versa*, just as usually assumed (von Fintel and Iatridou 2008: 119). Von Fintel and Iatridou (2008) focus on teleological modals, and only discuss deontic modals on the side. They suggest though that a similar picture might be applicable to deontic modals, where the secondary ordering source reflects less coercive norms than the primary one.⁶

Another proposal in the tradition of Kratzer is Alex Silk’s (2012), according to which weak necessity modals are conditional strong necessity modals. The idea is that ‘it should be that P ’ means roughly ‘if certain conditions C obtain, it must be that P ’, where the meaning of ‘must’ in the consequent is taken to be as suggested by Kratzer. The idea behind this account is that a characteristic feature of ‘should’ is that even if it should be that P , it is not the case that P has to obtain no matter what. For example, ‘you should keep your promises’ does

⁵This rests on the assumption that $f(w)$ is either circumstantial, i.e., reflects the geographical facts concerning Fife, or the epistemic modal background of a speaker who knows about these facts.

⁶Charlow (2013) endorses von Fintel and Iatridou’s proposal to introduce a second ordering source.

not mean that you have to keep your promises whatever may be. If saving one’s life demands breaking a promise, keeping the promise is not required, but ‘you should keep your promises’ nevertheless holds true. This analysis also has the advantage that it predicts that ‘must’ entails ‘should’, but not vice versa. If ‘you must do A ’ means that doing A is required under any circumstances, it of course entails that doing A is required under the specific circumstances C under which it is required according to the truth of ‘you should do A ’.

An alternative, second strategy is to deviate more strongly from Kratzer’s model. Finlay (2014: Ch.3; forthcoming) suggests to understand ‘should’ not as a universal quantifier akin to ‘all’, but rather as alike to ‘most’. Where ‘you must do A ’ means that you do A in all worlds of the relevant domain, ‘you should do A ’ means that you do so in most of them. Again, this captures that ‘must’ entails ‘should’ since ‘all’ entails ‘most’. Besides changing the modal force of the quantifier expressed by ‘should’ in contrast to ‘must’, Finlay only posits one conversational background, rather than two as Kratzer does. The modal domain this parameter creates consists of a modal background $\cap f(w)$, updated with a contextually salient goal or end z . ‘It should be that P ’ means on this account that in most worlds in the modal background in which the goal z is achieved, P rather than any of a set of contextually salient alternatives to P is the case (Finlay 2014: 73). ‘It should be that P ’ is accordingly the case iff it is more likely that the relevant goal z is achieved if P rather than any of certain alternatives to P occurs. This captures the thought that ‘should’ is a bit like ‘best’, in the sense that it describes what is the best means to achieve a particular end.

I would like to point out that all these semantic accounts for ‘should’ have in common that ‘should’ is context-sensitive according to them. This holds of the Kratzer-style inspired accounts by von Stechow and Iatridou as well as by Charlow since the conversational backgrounds (three in their case) are determined by context. On Silk’s account, ‘should’ is context-sensitive since the meaning of ‘should’ is, *inter alia*, composed of the meaning of ‘must’, which is context-sensitive, and because the specific condition C , under which one is required to do A if ‘You should do A ’ holds, is also contextually provided. Finally, the goal z that is part of the truth-conditions of should-sentences according to Finlay is also given by the context.

On all these semantics, context does not only fix whether ‘should’ expresses

epistemic or priority modality, but also which kind of priority modality or, even more specific, deontic modality it expresses. For example, ‘should’ can express what would be best or ideal in light of legal rules, the rules of morality or the rules of tennis. Depending on the particular theory, this is either a function of the ordering sources or of the ends that are contextually provided.

2.5 Information-Sensitivity

After having discussed some prominent semantic theories of the deontic ‘should’ (more or less) within the Kratzerian tradition, I will now turn to a characteristic of the deontic ‘should’ that is at the centre of this thesis: its information-sensitivity. I first explain what information-sensitivity is and then present how semantic theories in the literature account for it.

2.5.1 The Phenomenon

An information-state is usually considered to be a set of worlds (Kolodny and MacFarlane 2010: 130). Alternatively, information-states could be modelled as probability distributions over propositions, as suggested by, for example, Yalcin (2011: 299) and Wedgwood (forthcoming: sec. 3). As I use the term ‘information-state’, an information-states can represent a subject’s or a collective’s body of evidence, but can also just represent all facts, which might not be anybody’s evidence (or only God’s).

The idea that the deontic ‘should’ is information-sensitive amounts on a contextualist treatment of this phenomenon to the claim that sentences of the form ‘ S should do A ’ roughly mean ‘in light of information-state i , S should do A ’, where i is provided by the context in which the sentence is uttered. That is, there is no such thing as what one should do *simpliciter*, but only what one should do in light of an information-state. Furthermore, there must be at least two information-states i_1 and i_2 which are such that at least one instance of ‘ S should do A ’ is true in light of i_1 , but false in light of i_2 (Kolodny and MacFarlane 2010: 133). Finally, for the deontic ‘should’ to be fully information-sensitive, there also have to be at least two contexts that provide these two information-states i_1 and i_2 . Otherwise, it would only be possible, but never actually occur, that an instance of ‘ S should do A ’ is true if uttered in one and false if uttered in another

context of utterance. In the following, ‘information-sensitivity’ always means full information-sensitivity.

Before moving on to particular instances of information-sensitivity, it needs to be pointed out that the phenomenon need not be modelled by contextualist semantics. Kolodny and MacFarlane (2010) and in particular MacFarlane (2014: ch. 11) propose a relativist semantics for the deontic ‘should’, which captures its information-sensitivity as follows. Sentences of the form ‘*S* should do *A*’ are not true *simpliciter*, but only relative to so-called ‘contexts of assessment’. A context of assessment is, as the name suggests, a context from which a sentence is truth-evaluated, which is often not the same context as the one in which it was uttered, as when I evaluate a sentence you uttered. One of the contextual parameters of contexts of assessments are information-states. ‘*S* should do *A*’ can hence be true relative to one context of assessment, but false relative to another.

The main differences between contextualist and relativist semantics of the deontic ‘should’ with respect to information-sensitivity are accordingly the following two: First, contextualism entails that ‘*S* should do *A*’ expresses different propositions when uttered at two contexts with two different information-states i_1 and i_2 . In the first context, it means ‘*S* should do *A* in light of i_1 ’ and in the second ‘*S* should do *A* in light of i_2 ’. According to relativism, the sentence always expresses the same proposition. Second, for contextualists, if ‘*S* should do *A*’ is uttered in a context, it is true or false *simpliciter*. For relativists, it is only true or false relative to contexts of assessment. In the remainder of this chapter, and in fact in chapters 3 to 6 as well, I will only discuss information-sensitivity through the lense of contextualism. I will return to relativism in chapter 7. Besides relativism, a dynamic semantics of the deontic ‘should’ can also give an account of its information-sensitivity (Willer forthcoming: sec. 3.2). As explained above, I ignore dynamic semantics in this thesis for reasons of scope.

Let’s move on to particular instances of information-sensitivity. The distinction between the subjective and the objective ‘should’, as it is often made in ethics⁷ and also by some authors in epistemology⁸, can be seen as one. What a subject *S* should do in the objective sense of ‘should’ is what *S* should do in light of

⁷See Ross (1939), Prichard (1932), Ewing (1953), Brandt (1963), Jackson (1986), and Parfit (2011).

⁸See Brandt (1967), Gibbard (2005), and Wedgwood (forthcoming).

the relevant facts, and what *S* should do in the subjective sense, is what *S* should do in light of their evidence. The subjective and the objective ‘should’ can accordingly be seen as results of relativizing the deontic ‘should’ to different information-states.⁹

In the recent literature on the semantics for the deontic ‘should’, it has been pointed out that we also sometimes talk about what others should do in light of our collective evidence, i.e., their and our joint evidence.¹⁰ Such occasions seem to arise when we’re giving advice. Here is a much discussed example:

Miners.¹¹ Ten miners are trapped either in shaft 1 or in shaft 2. Floodwaters threaten to flood the shafts. Sean has enough sandbags to block one shaft, but not both. If Sean blocks one shaft, all the water will go into the other shaft, killing any miners inside it. If Sean blocks neither shaft, both shafts will fill halfway with water, and just one miner, the lowest in the shaft where the miners are, will be killed. Sean does not know in which shaft the miners are. He says:

(8) I should block neither shaft.

A physicist, who knows that the miners are in shaft 1, hears this and says to Sean:

(9) No, you should block shaft 1.

Intuitively, (8) and (9) are both correct. One way to explain this is that (8) is true since in light of Sean’s evidence, he should block neither shaft, and that (9) is true since in light of Sean’s and the physicist’s collective evidence, Sean should block shaft 1. Since the physicist is giving advice, she’s not talking about what Sean should do in light of just his evidence, but rather what he should do in light of their combined evidence, which includes her better-informed evidence.¹²

⁹For more on this, see chapter 3.

¹⁰See Kolodny and MacFarlane (2010), Björnsson and Finlay (2010), and Dowell (2013).

¹¹This case is introduced to the current debate about the information-sensitivity of deontic modals by Kolodny and MacFarlane (2010). They credit Parfit (ms) and Regan (1980).

¹²In chapter 5, I will argue that in advice cases like this, the advisee, here Sean, is actually also talking to the collective evidence, and that (8) is false. This does not undermine that there is a subjective ‘should’ as there are other cases where speakers talk relative to the subject’s evidence. See, for example, **Miners-CCTV** in subsection 2.5.3 and **Nostalgia** in subsection 2.6.2 below.

If this is the accurate explanation of the meaning of (9), **Miners** is a case in which ‘should’ is neither subjective nor objective, but relativized to the collective evidence. Contextualism brings order to this “profusion” (Jackson 1991: 471) of ‘should’s. It models them as different instances of one lexical entry for ‘should’, which result from ‘should’ being relativized to different contextually provided information-states.

Many authors argue that a specific theory capturing the information-sensitivity of ‘should’ is needed since Kratzer’s standard account of modals, in its original version, cannot do this.¹³ The ordering source g is realistic in that it ranks worlds in light of the facts, not in light of some agent’s less-than-omniscient information-state. For example, the standard Kratzerian ordering source in the context of Sean’s utterance ranks the worlds in the modal background relative to an ordering background containing the proposition $\langle \text{All miners are saved} \rangle$ ^{14,15} Worlds where Sean blocks the shaft in which the miners are are closer to this ideal than worlds where Sean blocks neither shaft. Thus, it is not true in all the best-ranked worlds that Sean blocks neither shaft. It follows that Kratzer’s theory can’t account for the truth of (8).

One might think that one could model information-sensitivity in Kratzer’s semantics by assuming that the modal base of an information-sensitive ‘should’ is epistemic, not circumstantial. That is, rather than containing all the worlds compatible with a certain realistic, i.e., true, description of how the world is in relevant features (e.g., that the miners are in shaft 1), it contains all the worlds compatible with what the speaker knows. This way, we can get the subject’s information-state to play a role in the truth-conditions of a sentence containing ‘should’, at least where the subject is also the speaker, as in the case of (8). However, this won’t get the results we need. The set of worlds compatible with what Sean knows contain both worlds where the miners are in shaft 1 and where they are in shaft 2. The worlds where the miners are in shaft 1 and Sean blocks shaft 1 and those where they are in shaft 2 and Sean blocks shaft 2 will be higher

¹³See Dowell (2013: 150), Cariani et al. (2013: 244), and Charlow (2013: 2303). For an objection to this, see Silk (2014: 709).

¹⁴Throughout this thesis, I use angle brackets to mark propositions.

¹⁵I’m ignoring here the additional complexities arising from modifying Kratzer’s theory of ‘must’ so that it works for ‘should’. Even when they are accounted for, the point about Kratzerian ordering sources being realistic stands.

ranked than those where he blocks neither shaft, since in the former kinds of worlds he saves 10 miners, where in the latter only 9. Thus, making the modal base epistemic rather than circumstantial won't help (Cariani et al. 2013: 235).

2.5.2 Option 1: Information-Sensitive Ordering Sources

Dowell (2012, 2013) suggests supplementing Kratzer's semantics by making the ordering source of the subjective 'should' dependent on information-states. In cases like (8), where a should-sentence is true since it is true in light of a subject's information-state, but not in light of the facts, the ordering source ranks worlds in light of the subject's evidence, not in light of the facts. An ordering source sensitive to Sean's lack of knowledge about the whereabouts of the miners is thought to rank a world where he blocks neither shaft higher than one where he blocks one of the shafts, even if the miners happen to be in that shaft. On this picture, the subjective 'should' is contextually sensitive to information-states in the sense that it is contextually sensitive to ordering sources, which again are sensitive to contextually salient information-states.

Before I move on to other suggestions of how to alter Kratzer's semantics to accommodate information-sensitivity, a note on terminology. Dowell (2013: 158) draws the distinction between the subjective and objective deontic 'should' such that the hallmark of the subjective 'should' is that it is information-sensitive, whereas the objective 'should' isn't since its ordering source is realistic. On my terminology, and the one of Kolodny and MacFarlane (2010: 117) and Wedgwood (forthcoming: sec. 4), on the other hand, the deontic 'should' is in general information-sensitive and the subjective and the objective 'should' are both instances of this information-sensitivity. The objective 'should' is also relativized to an information-state, but one that represents all the facts. I call this the *fully realistic information-state*. This might not be anybody's information-state (or only God's), but since information-sensitivity is a merely technical term, I don't see why we shouldn't construe the objective 'should' as a limiting case of information-sensitivity.

Dowell does not provide any motivation for her terminology, except that Angelika Kratzer suggested this to her in personal communication (Dowell 2013: 158, n. 22), and I don't think that there is some substantial disagreement here.

Dowell defines the objective ‘should’ as having an ordering source that is insensitive to somebody’s information-state, which is the same as having a ordering source that is sensitive to the fully realistic information-state. Thus, this all comes down to what we understand as an information-state. Following Kolodny and MacFarlane (2010: 117) and Wedgwood (forthcoming: sec. 4), I use the term in a more flexible fashion.

A further remark on terminology: when I speak of the *subjective* ‘should’, I only mean cases where ‘should’ is relativized to the subject’s information-state, unlike Dowell, who calls any instance of the deontic ‘should’ that is relativized to an information-state that represents somebody’s evidence, be it the subject’s, the speaker’s or the collective evidence, ‘subjective’. Where ‘should’ is relativized to, for example, the collective evidence, I call it *collective*. Again, I think this is a merely terminological issue and that nothing major hinges on this.¹⁶

Jennifer Carr (2015) argues that Dowell’s twist on Kratzer’s semantics has the problem that it does not get the following sentence right:

(10) If they are in shaft 1, I should block shaft 1.

Our intuition is supposed to be that (10) is correct. But imagine we are in a context in which Sean’s information-state is salient. In this context, (10) will turn out wrong. On Kratzer’s standard account of conditionals, the ‘if’-clause of an indicative conditional restricts the modal base of the modal in the consequent. The truth-conditions for an indicative conditional with a necessity modal in the consequent read on this picture as follows (Kratzer 1991: 648):

If $\llbracket \text{If } \phi, \text{ should } \psi \rrbracket^{w,c,f,g} = 1$ iff $\llbracket \text{should } \psi \rrbracket^{w,c,f^+,g} = 1$, where for any world v , w , $f^+(w) = f(w) \cup \{ v : \llbracket \phi \rrbracket^{c,f,g,v} = 1 \}$.

The expression with the curly brackets represents the set of worlds in which ‘ ϕ ’, the antecedent, is true. Thus, (10) is true at w iff Sean is blocking shaft 1 in all the $g(w)$ -best worlds in the set that contains all and only those worlds in the modal background $\cap f(w)$ in which the miners are in shaft 1. But this is not the case since ‘should’ is subjective and the ordering source g therefore ranks worlds

¹⁶For more on the subjective and objective ‘should’, see chapter 3. For more on the collective ‘should’, see chapters 5 and 6.

in which Sean blocks neither shaft still higher than those where Sean blocks shaft 1. One might argue that whenever (10) is uttered, the context changes such that ‘should’ becomes objective. In light of a realistic ordering, worlds where Sean blocks shaft 1 are ranked best. Carr replies that, intuitively, there is no change in the contextually salient priorities, i.e., in the ordering source, in the following discourse where (8) precedes (10):

(8) I should block neither shaft.

(10) If they are in shaft 1, I should block shaft 1.

Silk (2014) defends a proposal along the lines of Dowell, even though in a formally far more detailed manner. He suggests that the ordering sources g does not only take a world of evaluation as an argument, but also a contextually provided information-state, which ranks worlds in light of it. The problem Carr describes for Dowell’s account is avoided by the assumption that an ‘if’-clause updates not only the modal background of the ‘should’ in the consequent, but also the information-state that g takes as an argument.¹⁷ It would restrict this information-state to worlds in which the miners are in shaft 1. An ordering source which is determined by this updated information-state will rank those worlds best where Sean blocks shaft 1, which will make (10) true (Silk 2014: 707).

Carr’s (2015) own proposal is to have three, instead of two parameters. The first parameter is an information-state, which is in her model a pair of a modal background, i.e., the intersection of the propositions the modal base is mapped on to, and a probability function that assigns probabilities to the worlds in this modal background. The second parameter is a value function, which maps worlds to numbers, i.e., evaluates them. The third parameter is a decision rule. Decision rules rank the worlds in the modal background based on the probability function and the value function. For example, if the decision rule is to maximize expected utility, it ranks those worlds in the modal background best which according to the probability function of the first parameter and the value function of the second have the highest expected utility. (8) is predicted to be true on this account if we assume that the decision rule that is provided by the context in combination with Sean’s information-state ranks worlds in which he blocks neither shaft highest.

¹⁷In cases where the modal background is epistemic, the information-state and the modal background will be identical.

This will, for example, be the case if the principle is to maximize expected utility. Since Sean has no clue whether the miners are in shaft 1 or 2, Sean’s probability distribution assigns the propositions <The miners are in shaft 1> and <The miners are in shaft 2> a probability of 0.5, respectively. Assuming that the value function assigns saving all miners a value of, let’s say, 10, and saving 9 a value of 9, blocking neither shaft has an expected utility of 9 and blocking shaft 1 and blocking shaft 2 respectively an expected utility of 5.

(10) is correct, too, because the ‘if’-clause updates the information-state such that (i) the modal background is restricted to worlds where the miners are in shaft 1 and (ii) the probability function Pr is updated to $Pr^* = Pr(\cdot \mid \text{The miners are in shaft 1})$. This updated information-state will assign a probability of 1 to the propositions that the miners are in shaft 1. If the decision rule is, again, to maximize expected utility, this will yield the result that blocking shaft 1 is ranked highest and (10) accordingly true because blocking shaft 1 has the highest expected utility out of all alternative actions relative to the probability function Pr^* .

Notice that Carr’s account avoids any kind of commitment to a specific decision rule. It is not part of her semantics that the decision rule in the context of (8) or (10) is, for example, maximization of expected utility. Rather, she argues that the reason (8) and (10) are true is because the contexts of their utterance provide decision rules that make (8) and (10) true. The decision rule does not have to be ‘Maximize expected utility!’ Which decision theoretic principle ‘should’ expresses in these contexts is on Carr’s picture a matter of metasemantics, i.e., of how contextual parameters are resolved, not of the semantics of ‘should’ itself. This will be of relevance in the next subsection when we look at alternatives to Carr’s (2015) account.

Despite these differences, Carr’s (2015) theory incorporates information-sensitivity in a fashion very similar to Dowell (2013) and Silk’s (2014). As Carr (2015: 705) points out, her decision rules are generalizations of Kratzer-like ordering source; they perform the same function of ordering worlds in the modal background. Like Dowell (2013) and Silk (2014), Carr (2015) makes this ordering sensitive to a contextually provided information-state. The three hence account for the information-sensitivity of ‘should’ in the same way: the contextually provided ordering source ranks in light of a contextually provided information-state.

The reader might have noticed that all of these three different theories gloss over the distinction between ‘must’ and ‘should’. They all work within Kratzer’s standard semantics for ‘must’ and ignore the additional technical steps necessary to account for ‘should’, such as adding a second ordering source (von Stechow and Iatridou ms, 2008), treating ‘should’-sentences as hidden conditionals (Silk 2012), or changing the modal force of the quantifier (Finlay 2014, forthcoming). The same is true of the theories we’ll turn to in the next subsection. I assume the authors do so for the sake of simplicity and because it is reasonably clear how their treatments of information-sensitivity can be applied to the more complex semantic accounts of ‘should’. For example, all these accounts of ‘should’ rely on ordering sources in some way or another, so the treatments of information-sensitivity discussed in the present subsection can be applied to these accounts by making the relevant ordering sources information-sensitive. In the remainder of this chapter, I will follow these authors’ lead: in the next section 2.6, where I will discuss the doxastic ‘should’, I will first explain how the semantic accounts for the deontic ‘should’ mentioned can be applied to the doxastic ‘must’ and ‘should’; but when I explain, in the next step, how the models of information-sensitivity discussed here and in the next subsection can be applied to the doxastic ‘should’, I will ignore the difference between ‘must’ and ‘should’. This will save us getting into a lot of irrelevant complications.

2.5.3 Option 2: DP Models

Cariani et al. (2013) deviate farther from Kratzer’s account by adding a third parameter, a decision problem. I therefore call their model and the very similar one by Charlow (2013) *decision problem models*—or short, *DP models*. A decision problem is a function from a world to a set of ‘choosable’ actions. Choosable actions are those actions that an agent can do knowingly. For example, if I am faced with the problem of whether to turn right or left in world w , this decision problem will map w onto the following set of propositions: $\{ \langle I \text{ turn left} \rangle, \langle I \text{ turn right} \rangle \}$. It is crucial that the modal background is epistemic, not circumstantial on Cariani et al.’s account. That is, the modal background is the set of worlds compatible with the speaker’s knowledge. The decision problem forms a partition on the modal background by dividing it into cells representing the actions that

are choosable for the agent. In the miners case, this is the decision set Δ_{miners} :

$$\Delta_{miners}: \{ \langle \text{Sean blocks shaft 1} \rangle, \langle \text{Sean blocks shaft 2} \rangle, \langle \text{Sean blocks neither shaft} \rangle \}.$$

It is important that the proposition $\langle \text{Sean saves all miners} \rangle$ is not an element of it. This is so since saving all miners is not an action that Sean can do knowingly, given his ignorance about where the miners are. Finally, Cariani et al. (2013: 231) represent the ordering background for the miners case as the set Γ_{miners} :

$$\Gamma_{miners}: \{ \langle \text{All miners are saved} \rangle, \langle \text{At least 9 miners are saved} \rangle, \langle \text{At least 8 miners are saved} \rangle, \dots, \langle \text{At least 1 miner is saved} \rangle \}$$

Γ_{miners} ranks those propositions in the decision set Δ_{miners} highest that entail the most of the propositions in Γ_{miners} . Consequently, the proposition $\langle \text{blocking neither shaft} \rangle$ is ranked highest since it entails all propositions in Γ_{miners} except for $\langle \text{All miners are saved} \rangle$. The other two propositions entail none of the propositions in Γ_{miners} . Remember that the modal background is epistemic and thus contains both worlds where the miners are in shaft 1 and worlds where they are in shaft 2. Thus, the cell $\langle \text{Sean blocks shaft 1} \rangle$ of the modal background contains both worlds where Sean blocks shaft 1 and the miners are in shaft 1 and worlds where he blocks shaft 1 and they are in shaft 2. Thus, $\langle \text{Sean blocks shaft 1} \rangle$ entails none of the propositions in Γ_{miners} since it also holds at worlds where no miners are saved. The same holds for $\langle \text{Sean blocks shaft 2} \rangle$. Therefore, $\langle \text{Sean blocks neither shaft} \rangle$ is ranked highest and (8) turns out true.¹⁸

(10) is also correct since its ‘if’-clause restricts the modal background of the ‘should’ in the consequent to worlds where the miners are in shaft 1. The propositions describing Sean’s choosable actions within this modal background thus all only contain worlds in which the miners are in shaft 1. The proposition $\langle \text{Sean blocks shaft 1} \rangle$, which is part of this restricted decision set, therefore entails all propositions in the ordering background. If Sean blocks shaft 1 and the

¹⁸Charlow (2013) pursues a similar strategy. Following von Fintel and Iatridou (ms, 2008)’s suggestion to add a second ordering source to the semantics for ‘should’, he claims that the second ordering source ranks actions in contexts where the ‘should’ is sensitive to an information-state on basis of whether they are choosable in Cariani et al.’s sense. Blocking neither shaft is then said to be the action which is the highest ranked if the first ordering source, which ranks according to how many lives are saved, and the second one are merged.

miners are in shaft 1, this is guaranteed to save all 10 miners, and thus at least 9 miners, etc. The proposition is accordingly ranked highest and (10) true.

The information-sensitivity of ‘should’ is in this theory captured by the fact that which actions are choosable depends (i) on the knowledge, i.e., the information-state, of the subject the respective should-sentence is about and (ii) the epistemic modal background.

Carr (2015) objects to this proposal by Cariani et al. (2013) (and the similar one by Charlow 2013) that it carries a commitment to a specific decision-theoretic principle. ‘Sean should block shaft 1’ will turn out false even if Sean knows that there is an objective chance of 97% that the miners are in shaft 1. In this case, he cannot knowingly perform the action of saving all 10 miners since for this he would need to know that they are in shaft 1, which he doesn’t. The action <Sean blocks shaft 1> contains some worlds in which the miners are in shaft 2 since his knowledge does not rule this possibility out, even though he considers it highly unlikely. Again, blocking neither shaft will thus be the highest ranked action since it entails the most propositions in the ordering background.

This consequence of Cariani et al.’s account expresses a commitment to the decision-theoretic principle Maximin, which requires choosing that action whose worst possible outcome is the best among the worst possible outcomes of all alternative actions. Blocking shaft 1 has in the described scenario the highest expected utility (0.97×10 lives saved) compared to the alternative options of blocking neither shaft (1.0×9 lives saved) and blocking shaft 1 (0.03×10 lives saved). However, the worst possible outcome of blocking shaft 1 or blocking shaft 2, which is that no miners are saved, is worse than the worst possible outcome of blocking neither shaft, which is that 9 miners are saved. So Cariani et al.’s semantics has a commitment built into it to the Maximin principle, rather than for example the principle of maximizing expected utility. As Carr argues, however, such normative issues shouldn’t be part of the semantics of the term ‘should’. In other words, which decision-theoretic principle is correct should not be a matter of the meaning of ‘should’.¹⁹

I think there is a further problem with DP models: they cannot account for cases where the speaker is speaking relative to an information-state that is not their own. Consider the following variant of **Miners**:

¹⁹For an argument against this claim see Charlow (forthcoming).

Miners-CCTV. Everything is as in **Miners**, except that the physicist is not at the mine, but watches Sean via CCTV. After Sean utters (8), the physicist says:

(11) Yes, Sean should block neither shaft. Given what he knows, this is the only responsible thing to do.

The first sentence of (11) rings true, since the physicist apparently means by it that Sean should block neither shaft in light of his evidence, even though she knows that the miners are in shaft 1. According to DP models, (11) is false though. To see this, let's remember what the decision set is:

Δ_{miners} : {<Sean blocks shaft 1>, <Sean blocks shaft 2>, <Sean blocks neither shaft>}.

As Δ_{miners} is a partition on the epistemic modal background, which is the set of worlds compatible with the speaker's, i.e., the physicist's, knowledge, <Sean blocks shaft 1> only contains worlds in which the miners are in shaft 1. It follows that, just as in **Miners**, <Sean blocks shaft 1> entails all propositions in the ordering background Γ_{miners} and that it is hence ranked the highest. It follows that (11) is false as <Sean blocks neither shaft> is not ranked highest.

In reply to this problem, a defender of a DP model could give up the assumption that the epistemic modal background must be the set of worlds compatible with the speaker's knowledge. For example, assume that when the physicist utters (11), the epistemic modal background is the set of worlds compatible with Sean's knowledge, rather than hers; the reason for this could be that the physicist is trying to talk about what Sean should do in light of his knowledge, rather than in light of hers. If the decision set Δ_{miners} is a partition on Sean's, rather than the physicist's, epistemic background, <Sean blocks shaft 1> will contain both worlds where the miners are in shaft 1 and where they are in shaft 2. As a result, <Sean blocks shaft 1> won't entail <All miners are saved> and thus won't be the highest ranked cell, but rather <Sean blocks neither shaft> will.

To sum up, we have seen that Cariani et al. (2013) and Charlow's (2013) DP models struggle to account for cases where a better-informed speaker talks about what an agent should do in light of their inferior evidence. While the just

suggested move might fix this for the practical ‘should’, I will argue in the next section that it does not for the doxastic ‘should’.

2.6 Semantics for the Doxastic ‘Should’

Ideally, I would like to stay neutral between all these theories. It does not matter much for my aim of investigating the implications of the information-sensitivity of the doxastic ‘should’ how this phenomenon is exactly semantically modelled. All I need for my purposes is that there is some way to do it. In this section, I want to show how (some of) the theories discussed in sections 2.4 and 2.5 can be applied to the doxastic ‘should’. Subsection 2.6.1 explains how within those theories ‘should’ comes to express epistemic normativity. Subsection 2.6.2 argues that Dowell (2013), Silk (2014) and Carr’s (2015) semantic models can account for information-sensitivity in the doxastic ‘should’, whereas Cariani et al. (2013) and Charlow’s (2013) have difficulties with this.

2.6.1 Doxastic Ordering Sources

First, we need to understand how ‘should’ comes to express epistemic normativity. The general idea that I suggest is straightforward: on Kratzer’s semantics for ‘must’, the flavour of modality that ‘must’ carries is partly determined by the ordering source that is provided by the context in which ‘must’ occurs. ‘ S must do A ’ occurring at c means that morality demands that S does A if the ordering source provided by c ranks worlds according to whether S acts in accordance with moral norms. Similarly, ‘ S must believe P ’ occurring at c means that S is epistemically required to believe P if c provides an ordering source that ranks worlds according to whether S ’s doxastic attitudes satisfy epistemic norms. In the introduction I have characterized epistemic norms as being about what one should believe in light of considerations pertaining to truth, knowledge, justification, evidence, cognitive reliability, and the like. For example, a doxastic ordering source could be such that it contains all and only worlds where the relevant subject S believes all and only true propositions. This ordering source would reflect the truth norm according to which one should believe a proposition P iff P . In the following, I will refer to any ordering sources that reflects an epistemic norm as *doxastic*.

It depends on the particulars of your favourite semantics for the deontic ‘should’ how the story about how the doxastic ‘must’ expresses epistemic normativity carries over to the doxastic ‘should’. The options mentioned in section 2.4 were:

- Two ordering sources (von Fintel and Iatridou ms, 2008): ‘ S should do A ’ is true at a context c iff S does A in all those $g(w)$ -best worlds in the modal background $\cap f(w)$ that are ranked best by the secondary ordering source h , where f , g , and h are provided by c .
- Conditional strong necessity modals (Silk 2012): ‘ S should do A ’ is true at a context iff ‘If specific condition C obtains, S must do A ’ is true at c .
- Most-quantifier (Finlay 2014, forthcoming): ‘ S should do A ’ is true at a context c iff S does A in most worlds in the modal background $\cap f(w)$ updated by the goal z , where f and z are provided by c .

All these can be used as semantic models for the doxastic ‘should’. I suggest that on von Fintel and Iatridou (ms, 2008)’s account, ‘ S should do A ’ uttered at a context c expresses epistemic normativity if the two ordering sources g and h provided by c rank worlds according to epistemic norms. This assumes that there are two different kinds of epistemic norms, more coercive ones, that are reflected by the primary ordering source, and less coercive ones, reflected by the secondary ordering source. If there is no such distinction among epistemic norms, then von Fintel and Iatridou’s (2008) account leaves two options. First, one could conclude that there is no real distinction between the doxastic ‘should’ and ‘must’ and that they basically mean the same. A second, alternative option is to say that the doxastic ‘should’ is not a deontic, but another kind of modal. Von Fintel and Iatridou don’t agree on whether the non-epistemic ‘should’ is a deontic modal anyway (von Fintel and Iatridou 2008: 128). One of them suggests that it is rather a teleological modal. If we apply this interpretation to the doxastic ‘should’, we have an alternative account of under which conditions ‘should’ expresses epistemic normativity: when the ordering sources rank worlds according to whether the subject’s doxastic attitudes in those worlds are what they need to be like in order for the subject to achieve a certain epistemic goal, e.g., true belief or knowledge. The primary ordering source simply ranks all those worlds best in which this epistemic goal is achieved. The secondary ordering

source ranks all those worlds best in which the goal is achieved in the best manner. There might be no conceivable notion of achieving an epistemic goal in a better or worse manner, in which case we would end up, again, with the conclusion that there is no distinction between what one, doxastically, should and what one, doxastically, must believe.

On Finlay's account, the update of the modal background by a goal z effectively works as a ranking of the worlds in the modal background by an ordering source which ranks worlds in light of whether z is achieved in these worlds. 'Should' is doxastic on such an account if the contextually salient goal is an epistemic goal, similar to the view suggested by the second possible reading of von Fintel and Iatridou (ms, 2008),

Finally, on Silk's account of 'should' as a conditional strong necessity modal, 'should' is doxastic if the ordering source determining the modal flavour of the 'must' that occurs in the consequent of the conditional stating the truth-conditions for should-sentences is doxastic.

If we ignore the details of the specific accounts, what remains is that 'should' is doxastic if the ordering source(s) or goal supplied by the context in which it is uttered reflect(s) an epistemic norm or norms. I want to stay neutral on what these epistemic norms are. I will suggest some candidates in section 2.7, but the simplistic semantics for the doxastic 'should' that I will present there is purposively designed such that it's flexible enough to accommodate as many different epistemic norms as possible. In the spirit of Carr's remark that it is not the job of semantics to settle normative debates, it seems wrong to build a certain kind of epistemic norm into the semantics for the doxastic 'should'.

Before moving on to the issue of information-sensitivity, let me shortly address a question that might have come up in this subsection. As von Fintel and Iatridou's (2008) and also Finlay's (2014, forthcoming) discussion of 'should' suggest, we could classify the doxastic 'should' as a *teleological* rather than a *deontic* modal. If it is a teleological modal, it expresses which doxastic attitude one has to adopt in order to achieve an epistemic goal or *telos*, for example to believe only true propositions or only propositions one knows. Two remarks. First, I'm hesitant to categorize the doxastic 'should' as a teleological modal since this would amount to the claim that it is a semantic truth that epistemic instrumentalism is correct. This is the view that epistemic norms are instrumental norms, which tell us what

the means are to achieve certain epistemic ends. As this view is very controversial in epistemology,²⁰ I think we should reject, again in the spirit of Carr’s (2015) objection to Cariani (2013) and Charlow (2013), a semantics that would commit us to a particular view on this issue. If we consider the doxastic ‘should’ to be deontic we leave this open, since the rules expressed by deontic modals can be either instrumental or not. Second, it does not matter too much for the purposes of this PhD thesis whether the doxastic ‘should’ is deontic or teleological. What matters is whether it is information-sensitive, and it can be information-sensitive as a teleological modal just as well as it can be as a deontic modal. We can relativize what should be done in order to achieve a goal z to different sets of evidence. What one should do in order to achieve the goal of saving as many miners as possible in light of ignorance about where the miners are differs from what one should do for this purpose in light of knowing that they are in shaft 1.

2.6.2 Information-Sensitivity of the Doxastic ‘Should’

Let’s now turn to information-sensitivity. The following two types of semantic accounts of the information-sensitivity of the deontic ‘should’ were discussed in section 2.5:

- Information-sensitive ordering source (Dowell 2013, Silk 2014; Carr 2015): Ordering sources rank in light of the contextually salient information-state.
- Decision problem (Cariani et al. 2013; Charlow 2013): The ordering source does not rank worlds, but propositions that are classes of a partition on the epistemic modal background and represent those actions the subject can knowingly do.

I will discuss these two types in turn and claim that the first can incorporate the information-sensitivity of the doxastic ‘should’, while the second can’t.

What it means for a doxastic ordering source to rank in light of an information-state will depend on the exact nature of this ordering source. Here is just one option: a plausible epistemic norm is that one should proportion one’s belief to one’s evidence. Making the ordering source reflecting this norm sensitive to a contextually provided information-state i can be done if the ordering source

²⁰For powerful objections to it, see Kelly (2003), Jenkins (2007), and Berker (2013).

ranks those worlds highest in which the subject S is holding those beliefs that are proportioned to i , be that S 's evidence or another body of evidence. Alternatively, an ordering source might reflect the knowledge norm according to which one should believe a proposition P iff one is in a position to know that P . An ordering source reflecting this norm can be made sensitive to a contextually provided information-state i by designing it such that it ranks those worlds best where all of the subject S 's beliefs are such that if S held those beliefs on basis of i , they would be knowledge. I show in section 3.2 how my semantics of the doxastic 'should', which is supposed to be a simplified version of Dowell (2013), Silk (2014) and Carr's (2015), can reflect these and other epistemic norms and make them information-sensitive.

While Dowell (2013), Silk (2014) and Carr's (2015) semantics of the deontic 'should' hence can account for the information-sensitivity of the doxastic 'should', Cariani et al. (2013) and Charlow (2013)'s decision problem models—or short, *DP models*—can't. They capture information-sensitivity in their semantics by ranking propositions that are cells of a partition on the epistemic modal background induced by a decision problem. These propositions represent the actions the subject can knowingly do. If we want to apply this model to the doxastic 'should' what needs to be ranked are not propositions representing actions, but propositions representing doxastic attitudes. So let the propositions to be ranked be such that they describe which doxastic attitudes the subject can adopt knowingly. Under normal circumstances, these will be just any doxastic attitudes the subject can adopt, as most of the time when we, let's say, suspend judgement on whether P , we are in a position to know that we do so.

In subsection 2.5.3, we have already seen that DP models have difficulties with cases where the speaker is speaking relative to an information-state that is inferior to their own. This is not only so with respect to the practical 'should', but also with respect to the doxastic 'should':

Nostalgia. Renaud looks back at a situation at time t , when his evidence clearly supported disbelieving Q . We both know now that this evidence was misleading, and that Q is true. I say to Renaud:

(12) Don't beat yourself up. Disbelieving Q is what you should have

done at the time.²¹

The second sentence of (12) rings true, since I apparently mean by it that Renaud should have disbelieved Q in light of the evidence he had at t . Just like (11) in **Miners-CCTV**, (12) is false though, according to DP models. To see this, let's start with the assumption that the doxastic ordering source is 'realistic' and reflects an adequacy norm, which ranks doxastic attitudes by how 'close' they get to the truth. The adequacy norm will induce an ordering background along the following lines:

$$\Gamma_{adequacy}: \{ \langle \text{Renaud has a true belief about whether } Q \rangle; \langle \text{Renaud does not have a false belief about whether } Q \rangle \}$$

The decision set will look like this:

$$\Delta_{nostalgia}: \{ \langle \text{Renaud believes } Q \rangle; \langle \text{Renaud suspends judgement about whether } Q \rangle; \langle \text{Renaud disbelieves } Q \rangle \}$$

From all the doxastic attitudes Renaud can knowingly adopt, believing Q will be the highest ranked by $\Gamma_{adequacy}$. If Q is true, $\langle \text{Renaud believes } Q \rangle$ entails both $\langle \text{Renaud has a true belief about whether } Q \rangle$ and $\langle \text{Renaud does not have a false belief about whether } Q \rangle$. In contrast, $\langle \text{Renaud suspends judgement about whether } Q \rangle$ only entails $\langle \text{Renaud does not have a false belief about whether } Q \rangle$ and $\langle \text{Renaud disbelieves } Q \rangle$ entails none. It is important to notice that these entailment relations only hold because the decision set is a partition on the epistemic modal background. Because I know Q , $\langle \text{Renaud believes } Q \rangle$ only contains worlds where Q is the case and hence entails $\langle \text{Renaud has a true belief about whether } Q \rangle$. If I did not know Q , $\langle \text{Renaud believes } Q \rangle$ would also contain $\neg Q$ -worlds and the entailment relation wouldn't hold.

One might think that using a different ordering source will solve the issue for DP models. A natural candidate would be an ordering source that reflects the evidential norm. Its ordering background would look roughly as follows:

$$\Gamma_{evidence}: \{ \langle \text{Renaud adopts the doxastic attitude supported by his evidence} \rangle; \langle \text{Renaud doesn't adopt the opposite of the doxastic attitude supported by his evidence} \rangle \}$$

²¹This example is adapted from MacFarlane (2014: 298).

This ordering source will make (12) true since $\langle \text{Renaud disbelieves } Q \rangle$ entails all propositions in Γ_{evidence} . But this just moves the bump in the rug. As I argue extensively in section 3.3, the doxastic ‘should’ can have an objective reading. Now, imagine that I want to talk about which doxastic attitude Renaud should have objectively, i.e., in light of the facts, and utter:

(13) Renaud should believe Q .

The problem is that Γ_{evidence} makes (13) false since $\langle \text{Renaud believes } Q \rangle$ won’t be highest ranked as it entails none of the propositions in Γ_{evidence} : remember that Renaud’s evidence at t supported disbelieving Q . One could of course suggest that the ordering source shifts with the contextually provided information-state. That is, where the speaker uses the objective doxastic ‘should’, the ordering source is Γ_{adequacy} and where they use the subjective ‘should’, Γ_{evidence} is the ordering source. Moreover, we would have to add further ordering sources to account for sensitivity to other information-states, as for example the collective evidence. While this is a viable proposal, it will simply collapse DP models into the kind of model Dowell (2013), Silk (2012) or Carr (2015) suggest: the ordering source has become information-sensitive since it changes with the contextually provided information-state. A defining feature of DP models was, however, that the ordering source is not information-sensitive.

As suggested in subsection 2.5.3, a defender of a DP model could give up the assumption that the epistemic modal background must be the set of worlds compatible with the speaker’s knowledge. For example, assume that when I utter (12) the epistemic modal background is the set of worlds compatible with Renaud’s knowledge at t , rather than mine; the reason for this could be that I’m trying to talk about what Renaud should have believed at t in light of what he knew at the time, rather than in light of what I know now. If the decision set $\Delta_{\text{nostalgia}}$ is a partition on Renaud’s, rather than my, epistemic background, $\langle \text{Renaud believes } Q \rangle$ will contain Q - and $\neg Q$ -worlds since Renaud neither knew Q nor $\neg Q$. As a result, $\langle \text{Renaud believes } Q \rangle$ won’t entail $\langle \text{Renaud has a true belief about whether } Q \rangle$ because in some of the worlds in $\langle \text{Renaud believes } Q \rangle$ Renaud believes falsely.

However, while this avoids that $\langle \text{Renaud believes } Q \rangle$ is ranked highest, it does not make (12) true. For this, we need $\langle \text{Renaud disbelieves } Q \rangle$ to be ranked

highest and thus to entail more propositions in $\Gamma_{adequacy}$ than the other available doxastic attitudes. But it doesn't. Since the epistemic background contains some Q -worlds, $\langle \text{Renaud disbelieves } Q \rangle$ does also not entail $\langle \text{Renaud has a true belief about whether } Q \rangle$. Since $\langle \text{Renaud believes } Q \rangle$ and $\langle \text{Renaud disbelieves } Q \rangle$ in fact don't entail any of the propositions in $\Gamma_{adequacy}$, $\langle \text{Renaud suspends judgement about whether } Q \rangle$ at least entails $\langle \text{Renaud does not have a false belief about whether } Q \rangle$, which makes it the highest ranked proposition in $\Delta_{nostalgia}$.

To sum up, Cariani et al. (2013) and Charlow (2013)'s models might be able to capture cases where a better-informed speaker talks about what an agent should *practically* do in light of the agent's inferior evidence if they give up on the assumption that the epistemic modal background is always the speaker's. However, this move will not help them with respect to the doxastic 'should'.

One way to go here is to say that since their semantics for the information-sensitivity of the practical 'should' is to be preferred (for whatever reasons) and we should assume uniformity in the semantic models for the doxastic and practical 'should' (after all, they are both just instances of the deontic 'should'), the difficulty of incorporating a sensitivity of the doxastic 'should' to a subject's inferior information-state into their models shows that the doxastic 'should' is not information-sensitive in this way.

I think this conclusion is premature. First, uniformity in our semantic models is certainly an important theoretical value, but can be outweighed by considerations of empirical adequacy, i.e., the goal to respect linguistic intuitions. Second, as we have seen, models according to which the doxastic ordering source is information-sensitive are not limited in the same way as DP models. **Nostalgia** thus gives us actually a reason to prefer these accounts over DP models since they get us both: uniformity and empirical adequacy. That the DP models might not be the best of the available options anyway is supported by Carr's objections to them, which I have discussed in section 2.5. To conclude, the incapacity of the DP models to account, to the full degree, for the information-sensitivity of the doxastic 'should' rather counts against them, than against the claim that the doxastic 'should' is information-sensitive to this degree.

2.7 A Simplified Semantic Account of the Doxastic ‘Should’

Let’s take stock: whereas Dowell (2013), Silk (2014) and Carr’s (2015) semantics for the deontic ‘should’ can capture the information-sensitivity of the doxastic ‘should’, Cariani et al. (2013) and Charlow’s (2013) struggle to do this. The important insight is that it’s not in conflict with the “general wisdom” of current semantic theory to assume that the doxastic ‘should’ is information-sensitive.

In the following chapters, I will present evidence for the assumption that the doxastic ‘should’ is sensitive to a variety of information-states and discuss how this assumption can be applied to normative epistemology. For this purpose, I will need an account of the truth-conditions of sentences involving the doxastic ‘should’ that can accommodate its information-sensitivity. I will work with the following one:

Doxastic Should For any context of utterance c that provides a doxastic ordering source, any subject S , doxastic attitude D , proposition P , and information-state i , ‘ S should have D to P ’ is true at c iff having D to P reflects accurately how much i supports P , where i is the information-state provided by c .

I will spend the remainder of this subsection explaining this account and replying to potential objections.

Let’s start with the notation. What an information-state is I have already explained above; I will have a few more comments on this below. It should be clear what contexts of utterance, propositions, and subjects are. By ‘doxastic attitudes’ I mean both the coarse-grained ones—belief, disbelief, and suspension of judgement—as well as credences. I take the attitudes of believing P and disbelieving $\neg P$ to be equivalent.

I suggest that **Doxastic Should** is compatible with all the semantic theories discussed in subsection 2.5.2, which allow for full information-sensitivity of the doxastic ‘should’. Remember that on these ‘ S should have D to P ’ is true at a context c that provides a doxastic ordering source g iff S has D to P in all worlds which are highest ranked by g in light of the information-state i provided by c . Whatever doxastic ordering source g a semanticist might suggest, **Doxastic**

Should is flexible enough to be specified in such a way that the condition that having D to P accurately reflects how much i supports P is equivalent to the condition that in all those worlds that are ranked highest by g in light of i , S has D to P . This is at least so if the doxastic ordering source g is plausible in the sense that it reflects a familiar epistemic norm. As I argue below, **Doxastic Should** is that flexible since we have freedom in determining what goes into the information-state i and what it means for D to P to accurately reflect how much i supports P . Therefore, **Doxastic Should** can accommodate a wide variety of familiar epistemic norms, i.e., can be specified such that these norms turn out true. It can hence account for any information-sensitive ordering source a theorist might suggest, given that the ordering source reflects one of these familiar epistemic norms.

However, **Doxastic Should** is only a *simplified* version of the semantic theories that allow for full information-sensitivity of the doxastic ‘should’. Where most of these theories (Silk 2014 and Carr’s 2015) provide us with a semantic value for the doxastic ‘should’ itself, **Doxastic Should** only provides us with the truth-conditions for sentences in which a doxastic ‘should’ occurs unembedded. Therefore, it does not tell us how to deduce the meaning of sentences in which the doxastic ‘should’ is embedded, for example under propositional attitude verbs like ‘said’ or ‘believes’ or under the conditional ‘if’. That is, **Doxastic Should** only gives us limited instructions about the compositionality of the doxastic ‘should’ and how its semantic value interacts with the semantic values of other terms in a sentence in composing the meaning of a sentence. Since we will almost only be considering unembedded occurrences of the doxastic ‘should’ in this thesis, this is no great loss.²² I’ve decided to go with such a simplistic semantics in order to avoid unnecessary technical complexities. It is important to keep in mind though that **Doxastic Should** is precisely this—a *simplified* version of the semantics that allow for the information-sensitivity of the doxastic ‘should’—and not an *alternative* to them. The reader therefore does not have to worry whether embeddings of the doxastic ‘should’ can be taken care of by an information-sensitive contextualist semantics. They can. The only reason I operate in this thesis with **Doxastic Should**, which cannot deal with embeddings, is that it

²²I shortly address embedded occurrences in chapter 7, but won’t need a formal semantic model of the doxastic ‘should’ for this discussion.

makes things easier.

As already indicated, the reason why **Doxastic Should** can incorporate a variety of epistemic norms is that it leaves us freedom with respect to (i) what goes into the information-state and (ii) how we define what it means that D to P ‘accurately reflects’ how much i supports P . For example, **Doxastic Should** can accommodate the view that the truth norm for belief is the only epistemic norm, i.e., that the doxastic ‘should’ is strictly objective.²³ On this view, i is in every context the fully realistic information-state, which represents all the facts. If information-states are sets of worlds, this is the singleton containing the actual world. Where P is true in the actual world, the information-state entails P and where P is false, it entails $\neg P$. If we stipulate that (fully) believing P most accurately reflects an information-state entailing P , **Doxastic Should** amounts on this view to:

Truth-TC. For any context of utterance c that provides a doxastic ordering source, any subject S , and proposition P , ‘ S should believe P ’ is true at c iff P is true.

Similarly, **Doxastic Should** can account for a ‘universal’ evidentialism, according to which the only epistemic norm is that subjects should adopt the doxastic attitude that is supported by their evidence.²⁴ On universal evidentialism, all contexts supply the subject’s evidence as the information-state.

The claim that I will defend though is *Doxastic Contextualism*. This is the view that contexts provide *different* kinds of information-states and hence that the doxastic ‘should’ is fully information-sensitive in the sense explained in subsection 2.5.1.

A worry one might have is that **Doxastic Should** favours evidentialist interpretations of the subjective doxastic ‘should’. An evidentialist interpretation of the subjective doxastic ‘should’ states that one should subjectively have D to P iff one’s evidence supports having D to P . If we identify a subject’s information-state with their evidence, **Doxastic Should** would indeed be committed to such an

²³I will come back to the truth norm in section 3.2.

²⁴Feldman (2000) defends the view that a subject S should have D to P iff S ’s evidence supports D to P . I don’t know whether Feldman thinks that this is the *only* epistemic norm. I will come back to evidentialism in section 3.2.

evidentialist interpretation. In this case, it would be incompatible with norms according to which what one (subjectively) should believe is not simply what one’s evidence supports. For example, the knowledge norm requires that we don’t believe what we don’t know, so it requires that we don’t believe P if P is false, even if believing P is supported by our (in such a case misleading) evidence. Furthermore, the reliabilist norm that one should (subjectively) believe P iff one’s belief that P was reliably formed also clashes with said evidentialist reading of the doxastic ‘should’.²⁵ The fact that we’re free in determining what goes into the information-state comes here to our rescue again. Since ‘information-state’ is a technical term, it need not mean ‘evidence’ and hence **Doxastic Should** is not committed to an evidentialist reading. As I will show in subsection 3.4.1, by stipulating that one’s information-state is identical to the propositions one is in a position to know or to facts about the reliability of one’s belief-forming processes, **Doxastic Should** can make the knowledge or reliabilist norm come out true, too. As the reader can see, **Doxastic Should** is at least flexible enough to accommodate the familiar truth, knowledge, evidentialist, and reliabilist norms for belief.

I haven’t yet explained what it means for an information-state i to support a target proposition P and how we measure the strength of this relation. One natural idea is that the strength is proportionate to the evidential probability of P on i , i.e., the conditional probability of P on i according to the ‘objective’ probability function (Williamson 2000: ch. 10.1). A problem with this idea is that as every logically true proposition has a probability of 1, and thus also the conditional probability of any logical truth on information-state is 1, any information-state would support any logical truth to the maximum amount. It would follow that for every P that is logically true, ‘I should believe P ’ would be true in every context, whatever the contextually provided information-state is. This is intuitively wrong. It does not only make false predictions about the truth-values of doxastic-should sentences, but also seems to give a wrong picture of evidential support itself. It sounds odd to say that everybody has evidence in support of all logical truths. Views that reduce evidential support to inferential relations, rather than evidential probability, meet a similar fate. The strongest possible inferential relation is entailment. Hence every logical truth would be fully

²⁵For more on these norms, see subsection 3.4.1.

evidentially supported by every true information-state, since every logical truths is entailed by any truth.²⁶

How to specify what evidential support is is a vexed problem that any epistemological theory that makes use of the notion faces. Feldman (1995) claims that the traditional view on this issue is that the evidential support relation is a brute relation, which cannot be defined in other terms.²⁷ I don't take a stance on whether the notion of evidential support can be reduced or not, and I will simply rely on an intuitive notion of evidential support in the following. We should not require more than this for the purposes of this thesis.

While I do not have to say much about what evidential support is, I need to make clear that just because i entailing P or P having a high probability on i is not sufficient for i to strongly support P , it is nonetheless often the case that i supporting P is constituted by such deductive or probabilistic relations between i and P . For example, imagine you have drawn 100 balls out of a bag and 90 of them were black. The proposition <The next ball you draw will be black> has a probability of 90% on this evidence. It seems that the evidence supports the target proposition strongly due to this probabilistic relation in this situation and that a high credence in the proposition, a credence of 0.9 to be precise, accurately reflects this support relation.

A note on terminology: For the sake of brevity, I will sometimes speak of a doxastic attitude to a proposition P being *licensed* by an information-state i if the doxastic attitude accurately reflects the support relation between i and P .

Given the focus on the notion of evidential support in **Doxastic Should**, one might be concerned that **Doxastic Should** can't account for pragmatic encroachment, the alleged phenomenon that a subject's epistemic status does not only depend on the subject's evidence, but also on the subject's practical stakes. For example, Hawthorne (2004) and Stanley (2005) argue that whether one knows a proposition P partly depends on how relevant it is for one's practical concerns

²⁶Notice that this problem with inferential or probabilistic accounts of evidential support is not motivated by an internalist requirement that an agent needs to be aware of a relation between a body of evidence e and a proposition P for e to support P . That is, the problem is not supposed to be that we're not aware of how our evidence supports all logical truths, even though it does. After all, the existence of an evidential relation between e and P and an agent's awareness of this are distinct issues. Rather, it seems simply absurd to say that all bodies of evidence are evidence for all logical truths.

²⁷He counts Lehrer and Cohen (1983) and Fumerton (1990) as proponents of the view. Kelly (2007) makes the same point.

whether P is true. Regarding justification, Fantl and McGrath (2002) claim that whether one’s evidence meets the threshold for being justified in believing P is dependent on how practically significant it is for one whether P is true. If one is convinced by their arguments, one is probably sympathetic to the view that there is pragmatic encroachment on what one should believe, too. **Doxastic Should** can allow for this view, for example, if the information-state is chosen in the above described way so that **Doxastic Should** makes the knowledge norm true. If we combine this with the assumption that there is pragmatic encroachment on knowledge, we get the result that there is also pragmatic encroachment on what one should believe. Alternatively, we can modify **Doxastic Should** such that ‘ S should have D to P ’ is true iff having D to P accurately reflects (i) how much the contextually provided information-state i supports P and (ii) how much is at stake for S with respect to whether P is true. This modified account allows it to be the case that ‘ S should believe P ’ is true if P has a modestly high evidential probability on i and, practically speaking, it doesn’t matter much for S whether P holds, but false if i evidentially supports P to the same degree, but a lot hinges, for S , on whether P is true. I will stay neutral here with respect to the question of whether there is pragmatic encroachment on whether one should believe something and hence on whether **Doxastic Should** needs to be modified in the described way. This question is orthogonal to my main claim that the doxastic ‘should’ is information-sensitive and that this can solve a number of problems in normative epistemology. Hence, I will put the issue of pragmatic encroachment aside in the following.

2.8 Metasemantics

Doxastic Should tells us how we get the truth-conditions of a doxastic-should sentence from its parts and the information-state provided by the context in which the sentence is uttered. Being therefore a (very limited) theory about the compositionality of the doxastic ‘should’, i.e., the rules by which its semantic value contributes to the semantic value of a linguistic expression it is a part of, it is hence a form of *descriptive* semantics. One of the questions that *foundational* semantics or metasemantics is concerned with, in contrast, is how contexts provide the contextual parameters that feed into the composition of the meaning of

linguistic expressions at these contexts.²⁸ In the following chapters, we will often be concerned with the question which information-state is provided by a particular context. One tool we will use to answer this question is our linguistic intuitions. If a sentence is intuitively true, and is so only if one particular information-state is contextually provided, it speaks in favour of this information-state being contextually provided. However, we do not always have clear linguistic intuitions, and the linguistic intuitions of two competent speakers can clash. Hence, we don't want to rely on them alone. Some theoretical considerations on the metasemantical question of how information-states are chosen would be very helpful.

Janice Dowell (2011, 2012, 2013) gives an account of the metasemantics of epistemic and deontic modals. As she points out, the issue of how context provides contextual parameters comes up with regard to other context-sensitive expressions, too, and a lot of research has been done on this. A much contested issue is, for example, which contextual feature determines the reference of demonstratives like 'this' or 'that'. On the one hand, speaker intentions seem to play an important role. Imagine I stand in front of a painting and say 'This is beautiful'. That I had the painting in mind and intended to refer to it by the demonstrative 'this' seems to explain why the semantic value of 'this' is the painting. On the other hand, intentions probably can't do all the work. Kaplan (1978: 239) gives a case in which he points behind him to a picture, which he thinks is of Carnap, but in fact shows Spiro Agnew (the 39th US Vice President), and says: "[That] is a picture of one of the greatest philosophers of the twentieth century". While he intended to refer to a picture of Carnap, it doesn't seem like he succeeded. As Kaplan (1989b) argues, what fixes the reference here might not be the speaker's intention, but rather his demonstration. If Kaplan is pointing towards a picture of Agnew, this picture is what 'that' refers to, and not the picture of Carnap he has in mind. This claim is supported by the consideration that what 'that' refers to is part of what is said by Kaplan's utterance. As the audience can't possibly know what Kaplan's intention is, the audience couldn't know what has been said if this intention fixed the reference of 'that'. But what is said is supposed to be what is communicated between discourse participants and Kaplan's audience should

²⁸The distinction between descriptive and foundational semantics is from Stalnaker (1997). (Kaplan, 1989a: 573f.) makes a similar distinction between semantics and metasemantics. I will speak of 'metasemantics' in the following, since it is the more commonly used term.

therefore be capable of recovering what has been said (Dowell 2011: 4).

The issues here are complex, and I don't want to get too deeply into them. What is worth pointing out though is that even if intentions fix the reference of demonstratives, speakers need to rely on certain extra-intentional cues in order to clearly communicate what they have in mind. Having changed his position in Kaplan (1989a), Kaplan proposes that speaker intentions *do* fix reference. However, demonstrations are needed to convey these intentions to the audience such that the audience can pick up on what has been said.²⁹

As Dowell (2011, 2013) points out, parallel issues arise with respect to the information-sensitivity of epistemic and deontic modals. As I have explained above, modals are treated as quantifiers in standard semantic theories. 'Ordinary' quantifiers, as in 'Every beer is in the bucket', rely on context to determine the modal domain they are quantifying over (Stanley and Gendler Szabó 2000). As in the case of demonstratives, speaker intentions seem to play a role here. Since information-states function as restrictions on the modal domain of deontic modals, it is natural to think that which information-state is provided by a context to do the restricting work is up to the speaker, just as it is up to the speaker to fix the domain of ordinary quantifiers. Remember the following case from subsection 2.6.2:

Nostalgia. Renaud looks back at a situation at time t , when his evidence clearly supported disbelieving Q . We both know now that this evidence was misleading, and that Q is true. I say to Renaud:

(12) Don't beat yourself up. Disbelieving Q is what you should have done at the time.

What I seem to mean by the second sentence of (12) is that Renaud should have disbelieved Q in light of the evidence he had at t . In the terms of **Doxastic Should**, the contextually provided information-state is Renaud's evidence at t . A plausible explanation for this being the context's information-state is that I *want* to speak relative to Renaud's evidence. The reason I intend to speak relative to his evidence is that I'm trying to evaluate whether he was epistemically rational

²⁹Perry (2009) develops Kaplan's (1989a) account further. For further discussion of the role of speaker intentions in fixing the semantic value of demonstratives, see, among others, Gauker (2008), Stokke (2010), and King (2014).

when he came to disbelieve Q , in other words, whether he had evidence for what he believed. But now consider the following variant of **Nostalgia**.

Progress. Renaud and I are looking back at the same situation at time t as in **Nostalgia**. We have just discussed what is required to be epistemically rational, which is to form one's doxastic attitudes on basis of the evidence one has at the time when one forms those attitudes. However, without telling him, I'm looking at Renaud's past situation not in terms of what would have been epistemically rational for him to believe, but what would have been the best for him to believe given what we know now. Thus, I'm intending to speak to our current evidence, when I say:

(14) Renaud, you should have believed Q .

What I mean by (14) is 'In light of our current evidence, Renaud should have believed Q .' But is this what I said? After all, Renaud couldn't tell. We've just been discussing epistemic rationality, so it's plausible for him to assume that what I intended was to speak relative to his evidence at the time (and came to the wrong conclusion about what this evidence supported).

The context in which I utter (14) violates what Dowell (2011: 1, 2013: 150) calls the "publicity constraint" on contexts. This constraint requires of a context that it is "capable of publicly manifesting a speaker's parameter-determining intentions" to a reasonable audience (Dowell 2013: 1). The relevant features of the context in which I utter (14) do not manifest my relevant intentions. Given that we've just been discussing epistemic rationality, epistemic rationality is conversationally salient. This is one of the context's features, and it indicates that I'm intending to speak relative to the evidence Renaud had at t .

It is not clear to me whether Dowell thinks that violating this constraint (i) leads to the context failing to provide an information-state *or* (ii) only leads to the context failing to signal to a reasonable audience which information-state has in fact been provided by the context in virtue of the speaker's (my) intention. In

any case, I don't need to commit to one of these two options here.³⁰

Besides conversational salience, the pragmatic purpose of an utterance can also indicate a speaker's intentions. That a speaker utters a doxastic-should sentence for the pragmatic purpose of evaluating your epistemic rationality is good evidence that they are intending to speak to your information-state, given what epistemic rationality requires. If they want to evaluate your epistemic rationality, they should talk about which doxastic attitude you should have in light of your evidence. Therefore, even if this is not a speaker's actual intention, it is reasonable for an audience to assume that the speaker is talking relative to the addressee's evidence if they are engaged in evaluating the addressee's rationality.

One might be convinced by the argument that contextual features must be recognizable to the audience in order for them to determine which parameter is provided by the context. In this case, one could argue that the publicity constraint is not only necessary for a context to provide a parameter—as according to option (i)—but also sufficient.³¹ On this account, if the pragmatic purpose of an utterance is clearly to evaluate an agent's epistemic rationality, but the speaker is mistaken about which information-state she should refer to and intends to refer to a wrong one, then the information-state provided is the one that is suited best for the goal of evaluating epistemic rationality, and not the one the speaker has in mind.

To sum up, we have three different contextual features that could determine the context's information-state or at least have a signalling function with respect to which the contextually provided information-state is: the speaker's intention, the utterance's pragmatic purpose and what's conversationally salient. What to do now? I won't commit to one of these features being the one that determines the context's information-state since the question of how context parameters are fixed is notoriously difficult and applies to many other context-sensitive terms. Rather, when judging what is the contextually provided information-state, I will consider all of these three contextual features.

³⁰That Dowell states that her account is inspired by Kaplan (1989a) indicates that she would favour option (ii). However, there are other formulations in Dowell (2011, 2013), which suggest the opposite. King (2014) defends a view akin to (i) with respect to demonstratives, according to which it is necessary for an object to be the semantic value of a demonstrative that a competent hearer can recognize that this object is the semantic value.

³¹Gauker (2008) defends such a view with respect to demonstratives.

2.9 Conclusion

This chapter introduced the current literature on the semantics for the deontic ‘should’. Since this literature focuses on the practical ‘should’, I had to show that the discussed theories also give us a semantics for the doxastic ‘should’. I argued that only some of them, Dowell (2013), Silk (2014) and Carr’s (2015), can account for information-sensitivity in the doxastic ‘should’, and that others, the very similar proposals of Cariani et al. (2013) and Charlow (2013), can’t. Furthermore, I presented a simplistic semantics for the doxastic ‘should’, compatible with Dowell (2013), Silk (2014) and Carr’s (2015) theories, according to which doxastic-should sentences are sensitive to contextually provided information-states. I coined the view that such sentences are relativized to different kinds of information-states in different contexts *Doxastic Contextualism* and I will provide evidence for it in the following chapters 3 to 5. Finally, I explained which contextual features can play a role in determining which information-state a context provides.

Chapter 3

The Subjective and the Objective Doxastic ‘Should’

3.1 Introduction

My aim in this chapter is to provide a first piece of evidence for Doxastic Contextualism by presenting a first instance of the information-sensitivity of the doxastic ‘should’: the distinction between the subjective and the objective doxastic ‘should’. A parallel distinction for the moral ‘should’ is well-known in ethics, but it has also been discussed in epistemology with respect to the doxastic ‘should’. In section 3.2, I show that this distinction can be integrated into the model for the information-sensitivity of the doxastic ‘should’ that I have introduced in chapter 2. While it is widely accepted that there is a sense of the doxastic ‘should’ that is subjective, there will be more resistance to the idea that there is an objective doxastic ‘should’. Therefore, I will present five arguments supporting the idea that there is such a ‘should’. Among other things, I argue that we use the objective doxastic ‘should’ in everyday discourse to describe what’s epistemically best (section 3.3).

In section 3.4, I explain how attention to this distinction can be of use to epistemologists. A number of authors have proposed a truth norm for belief according to which one should believe a proposition P iff P is true. This norm is apparently in conflict with other norms discussed in the literature, which make whether one should believe P not (only) dependent on whether P , but (also) on

one's epistemic situation with respect to P , for example whether one's evidence supports P or whether one knows P . In what follows, I will call such norms *subjective*. I claim that the 'should' in the formulation of the truth norm is objective, whereas the 'should' in the formulation of the other norms is subjective. One upshot of the information-sensitivity of the doxastic 'should' with respect to the debate about the truth norm and its apparent rivals is thus that these norms are actually not in conflict. In section 3.5 I reply to four possible objections. The first three all attack, in some way or other, the truth norm and hence suggest that the conflict should be resolved by giving up the truth norm. I will defend the truth norm against these objections. The fourth objection, by Gibbons (2013), criticizes that distinguishing between the subjective and objective doxastic 'should' ducks the question of which of the two senses of 'should' guides doxastic deliberation. I reply that I am happy to grant that the objective doxastic 'should' is not guiding. This fits with the point I made in section 3.3 that we use the objective 'should' in ordinary discourse to describe what's epistemically best.

3.2 The Subjective/Objective Distinction

The distinction between the subjective and objective 'should' is often used by those who work on ethics or theories of rationality, i.e., those who are concerned with what we practically should do. It is roughly the distinction between what a subject should do in light of their epistemic perspective and what a subject should do in light of the facts.¹ Imagine that you're in a hotel room and the hotel is on fire, but you're not aware of it. In some sense of 'should', you should jump out of the window (assuming there is no other safe exit). Since the hotel is on fire, this is the best available action for you. In some other sense of 'should', you should not jump. Given that you don't know that the hotel is on fire, it would be irrational of you to simply jump out of the window.² We seem to have a puzzle: on one hand, you should jump out of the window. On the other hand, you should not. Distinguishing between the subjective and objective 'should' solves the puzzle: subjectively, in light of your evidence, you should not jump out of the window.

¹See Ross (1939), Prichard (1932), Ewing (1953), Brandt (1963), Jackson (1986), Parfit (2011), Kolodny and MacFarlane (2010), and Wedgwood (forthcoming).

²This example is from Parfit (1997: 99).

Objectively, in light of the fact that the hotel is on fire, you should jump out of the window.

Drawing the same distinction in epistemology is less common, but some authors have done so.³ It is the analogous distinction between what one should believe in light of the facts and what one should believe in light of one's epistemic perspective.⁴

The facts that the objective sense of the practical 'should' is relativized to are often described as *all* the facts.⁵ I assume the same for the objective doxastic 'should'. Which doxastic attitude one objectively should adopt towards a specific proposition P depends only on one fact in the set of all true propositions. If this set contains P , i.e., if P is true, one objectively should believe P . If this set contains $\neg P$, i.e., if P is false, one objectively should believe $\neg P$ (or disbelieve P). We can give the following corresponding definition of the objective doxastic 'should':

Objective Doxastic Should For any subject S , proposition P , and time t , S should objectively believe P at t iff P is true.⁶

Of course, the 'should' in this definition is supposed to be doxastic. That is, it does not state that one should, for practical reasons, objectively believe P iff P . I claim that the objective doxastic 'should' occurs in the formulation of the truth-norm:

³See Brandt (1967), Gibbard (2005), and Wedgwood (2007: ch. 5.2).

⁴There is a quite widely discussed distinction between subjective and objective *justification*, which is not to be confused with the distinction that I have in mind here. Roughly, one is objectively justified to have a doxastic attitude D to a proposition P iff one is actually justified to have D to P ; and one is subjectively justified to have D to P iff one (justifiably) believes that one is objectively justified to have D to P . (See Goldman 1986: 73; Pollock 1979: 109f.; Alston 1985: 62; Kornblith 1985: 264; Moser 1985: 62; and Kvanvig 1984: 71.) Sometimes, this distinction is also described as a distinction between a subjective and objective 'should' in epistemology, where the first expresses what one should believe in light of the actual evidence or reasons one has, and the second expresses what one should believe in light of what one (justifiably) believes one's evidence or reasons to be (Pollock 1979: 109f.; Kornblith 1985: 264; and Moser 1985: 62). I do not want to deny that there is such a distinction between senses of 'justification' or 'should'. In fact, I think that this distinction can be accommodated by the model for the information-sensitivity of the doxastic 'should' that I have suggested in chapter 2. While I am accordingly not opposed to the distinction between subjective and objective justification, it is not the distinction between subjective and objective 'should' that I am after here.

⁵See Kolodny and MacFarlane (2010: 117). Similarly, Wedgwood (forthcoming: sec. 4) speaks of the objective 'should' as relativized to an "omniscient" information-state.

⁶If propositions are time-indexical, we need to add 'at t ' to the right of the biconditional.

Truth For any subject S , proposition P , and time t , S should believe P at t iff P is true.⁷

The subjective doxastic ‘should’ expresses, roughly, what one should believe in light of one’s evidence. I suggest the following definition:

Subjective Doxastic Should For any subject S , proposition P , doxastic attitude D and time t , S should subjectively adopt doxastic attitude D to P at t iff having D to P reflects accurately how much S ’s information-state at t supports P .

I have chosen the technical term ‘information-state’ instead of evidence on purpose. This allows **Subjective Doxastic Should** to be open to different accounts of the subjective doxastic ‘should’. In subsection 3.4.1, I will argue that due to this flexibility many epistemic norms in the literature that I call *subjective* can be interpreted as being formulated in terms of the doxastic subjective ‘should’. A subjective epistemic norm makes what a subject should believe dependent on the subject’s epistemic perspective. A good example is the following norm:

Evidentialism For any subject S , proposition P , doxastic attitude D and time t , S should have D to P at t iff D to P is supported by S ’s evidence at t .⁸

It is easy to see that **Evidentialism** comes out true on a subjective reading of ‘should’ if we assume that the subject’s ‘information-state’ that is mentioned in **Subjective Doxastic Should** is the subject’s evidence.

The subjective/objective distinction can easily be accommodated by the model for the information-sensitivity of the doxastic ‘should’ that I presented in chapter 2:

Doxastic Should For any context of utterance c that provides a doxastic ordering source, ‘ S should have doxastic attitude D to P ’ is true at c iff having D

⁷I will discuss the norm in more detail in sections 3.4 and 3.5.

⁸See Feldman (2000: 678). For the sake of brevity, I will drop the quantification over propositions, subjects, doxastic attitudes, and times and the reference to times in most of my formulation of norms in the following. Unless otherwise stated, the reader should take them to be implicit.

to P reflects accurately how much i supports P , where i is the information-state provided by c .⁹

The doxastic ‘should’ takes on a subjective sense where the contextually provided information-state i is S ’s information-state (whatever that exactly is) at the time t of which the speaker says that S should have D to P then. It takes on an objective sense where i is all the facts. As mentioned in subsection 2.5.2, I will call this latter information-state the *fully realistic information-state*. If we look at the objective doxastic ‘should’ in this way, it is a limiting case of the information-sensitivity of the doxastic ‘should’.¹⁰

On this picture, the doxastic ‘should’ is strictly speaking not ambiguous between two senses. Rather, it is akin to an indexical like ‘I’. The sentence “I am hungry” has different truth-conditions in different contexts, i.e., when uttered by different speakers, but we wouldn’t say that ‘I’ has different senses. To put it in Kaplan’s (1989a, 1989b) standard terminology for indexicals, ‘I’ has different *contents* in different contexts, i.e., refers to different persons, but only one *character*. That is, there is a unique rule according to which ‘I’ is mapped in a context to a content. Analogously, ‘should’ has different contents in different contexts and can, for example, mean ‘should in light of the subject’s evidence’ in one context and ‘should in light of the facts’ in another. However, it has one character, i.e., there is one rule determining which content ‘should’ has in a context.

What unites context-sensitive and ambiguous terms is that both kinds of terms can have multiple meanings. What distinguishes them is that context-sensitivity is characterized “by interaction with (extra-linguistic) context” (Sennet 2016), whereas ambiguity is not. That is, which specific meaning (content, in Kaplan’s terms) a context-sensitive term has is a function of the context in which the term occurs. *Strictly* ambiguous terms can easily be distinguished from context-sensitive ones since the former have multiple meanings that are unrelated, as in the case of ‘bark’, which is ambiguous between the sound dogs make and the outermost layer of trees.¹¹ It is harder to distinguish context-sensitive and polysemous terms. The

⁹For the sake of brevity, I have dropped the quantification over subjects, doxastic attitudes, propositions, and information-states in my formulation of the truth-conditions of doxastic-should sentences here and in the following, except where I deem it helpful to make it explicit.

¹⁰As I also mention in subsection 2.5.2, Dowell (2013) construes the objective ‘should’ as information-*insensitive*. As I explain there, her position is only superficially at odds with mine.

¹¹See Gross (2001: 9).

multiple meanings of polysemous terms are somehow related, for example in the form of a type-token relation, as in the following case:

- (1) He left the bank 5 minutes ago. He left the bank 5 years ago.

‘Bank’ in the first sentence refers to a specific building, whereas in the second sentence it refers to an institution that owns this building.

It is clear that the subjective/objective distinction is not a form of strict ambiguity. Furthermore, what speaks in favour of it being context-sensitive rather than polysemous is that we can give a semantic model of the doxastic ‘should’—like the ones discussed in section 2.6—that gives us a unique rule that fixes a contextual parameter, i.e., an information-state, to the subjective and objective meanings of the doxastic ‘should’. Thus, we can give a plausible semantic treatment of the doxastic ‘should’ on which the subjective/objective distinction is the result of interaction with linguistic context.

I could go into a deeper linguistic analysis here to settle the issue whether the subjective/objective distinction is a form of ambiguity or context-sensitivity. I don’t think that this is necessary though. What matters to the discussion in this chapter is that ‘should’ can carry a subjective and objective meaning. Whether this is so in the form of different Kaplanian contents due to context-sensitivity or in the form of different senses due to ambiguity is not important. Furthermore, it is also not decisive for the general project of my thesis. My claim is that the doxastic ‘should’ is relativized to different information-states and that this fact has interesting implications for normative epistemology. Whether the relativizability to information-states is a form of ambiguity or context-sensitivity is, in the light of the purposes of the project, a mere technicality.¹²

3.3 The Objective Doxastic ‘Should’

In this section, I want to present positive arguments in favour of my claim that there is an objective doxastic ‘should’. This won’t be all I have to say in support of the objective doxastic ‘should’. I will defend the truth norm against common

¹²Even though I favour the view that the subjective/objective distinction is a form of context-sensitivity, I will continue using the term ‘sense’ for the sake of simplicity.

objections to it in section 3.5, which amounts to a defence of the objective doxastic ‘should’.

There are few who would deny that there is a subjective doxastic ‘should’ in the sense intended here.¹³ As Gibbons (2013: 33) points out, epistemologists think of norms as mostly subjective, as being determined by the subject’s perspective and not by external facts alone. Even on the most ‘objective’ of all subjective norms, the knowledge norm, which makes whether one should believe a proposition P sensitive to the external fact of whether P is true, one’s epistemic perspective, namely, whether one is justified, plays a role in determining whether one should believe P .¹⁴

That there is a subjective doxastic ‘should’ is also supported by linguistic evidence. Consider this case from the previous chapter:

Nostalgia. Renaud looks back at a situation at time t , when his evidence clearly supported disbelieving Q . We both know now that this evidence was misleading, and that Q is true. I say to Renaud:

- (2) Don’t beat yourself up. Disbelieving Q is what you should have done at the time.

The second sentence in (2) rings true. This indicates that there is a doxastic ‘should’ in English that is relativized to the subject’s information-state, and not the facts.

If you’re a fan of something akin to the knowledge norm of belief, the following example might be more to your taste. Imagine that you recklessly believe a proposition P without having evidence for it, which (luckily for you) is actually true. It seems appropriate for me, who knows P , to say to you:

- (3) You should not believe P ; you don’t know P !

If the doxastic ‘should’ only had an objective, but not a subjective sense, then (3) would be false. However, (3) sounds true.

¹³True, there are those who think that there are no epistemic norms. I’ve addressed this position in the introductory chapter 1 and will ignore it here.

¹⁴In the following, I will often use the formulation ‘what one should believe’ as shorthand for the more complex formulation ‘which doxastic attitude one should adopt’. This is *not* supposed to mean that what is said only applies to the coarse-grained attitude of belief but not to other doxastic attitudes such as credences.

Let's turn to the objective 'should'. I suspect that here we might face more resistance than in the case of the subjective 'should'. The flip side of the fact that epistemologists think about normativity mostly from a subjective standpoint is that the notion of what one objectively should believe might sound strange to many. Feldman (1988) discusses Brandt (1967), who was, to my knowledge, the first who drew the distinction between what one should believe subjectively and what one should believe objectively, i.e., in light of the facts. Feldman (1988: 409) claims that this distinction is actually alien to epistemology and that Brandt "mentions no one who does draw it in this way and I [Feldman] have found no one who does so." Feldman then goes on to argue that the subjective/objective distinction that matters in philosophy is the one between subjective and objective justification, i.e., a distinction between different notions of what one *subjectively* should believe.¹⁵

The claim that the doxastic 'should' has an objective sense is in need of support. I will thus present five arguments in support of it in the following. I don't expect that every reader will be convinced by every argument, but if you just accept one of them, you have compelling reason to believe that the doxastic 'should' can be objective.

First, an argument that is provided in favour of the truth norm proceeds from the following claim:

Correct Belief A belief that P is correct iff P is true.

It is then argued that 'correct' is a normative term and that we can therefore infer **Truth** from **Correct Belief**.¹⁶ **Truth** can only be true if the 'should' in its formulation is objective. Thus, if the argument from **Correct Belief** to **Truth** is sound, the doxastic 'should' has an objective sense.

Second, there is linguistic evidence suggesting that we at times employ an objective 'should'. Consider the following scenario:¹⁷

Police-CCTV-1. Detective Lester Freamon is pondering about who killed Bob. He knows that it was either Omar or Marlo. Furthermore,

¹⁵For more on this see footnote 4 above.

¹⁶See Gibbard (2005) and Engel (2007).

¹⁷This case, and variants of it that I will discuss in this and following chapters, is structurally analogous to the **Miners** case, which I introduced in subsection 2.5.1.

his reliable informant Bubbles told him that Marlo was at the other end of town at the time of the murder. He therefore concludes that it was Omar. His colleague Kima Greggs, who is watching Lester via CCTV, knows that Bubbles was mistaken this time and that Marlo in fact killed Bob. She utters:

(4) Oh what a shame! Lester should believe that it was Marlo.

(4) rings true, and it is true if the ‘should’ occurring in it is objective. In light of the fact that Marlo killed Bob, Lester should believe that Marlo killed Bob. (4) can’t be true if the ‘should’ in it is subjective as Lester’s epistemic perspective supports believing that it was Omar, not Marlo. Thus, (4) is a sentence in which the objective doxastic ‘should’ occurs.

A potential objection to this line of reasoning is to argue that the ‘should’ is neither subjective nor objective, but relativized to another information-state. In particular, I propose myself in chapter 5 that we sometimes employ a collective sense of the doxastic ‘should’, where the information-state is the joined evidence of the speaker and the subject. On such a collective reading, (4) would be true as well since in light of Lester’s and Kima’s joint evidence, Lester should believe that it was Marlo.

I reply that the purpose for which we employ ‘should’ in (4) is different from the use to which we put the collective doxastic ‘should’. As I explain in chapter 5, the collective doxastic ‘should’ is used when we give others advice on which doxastic attitude to adopt. Kima can’t be said to give Lester advice as Lester can’t hear what Kima is saying. Rather, what Kima is doing by uttering (4) is assessing Lester’s belief in light of the standard of correctness for beliefs. A correct belief is true, and that’s why she’s expressing (mild) frustration with the fact that Lester believes something false (‘What a shame!’) and claims that he would better believe that Omar killed Bob.

Further evidence that the sense of the doxastic ‘should’ that we employ when we evaluate someone’s belief is not the same as the one we use for the purpose of epistemic advice is the following variant of **Police-CCTV-1**:

Police-CCTV-2. Lester is pondering about who killed Bob. He knows that it was either Omar or Marlo. Furthermore, his reliable informant Bubbles told him that Marlo was at the other end of town at the time of the murder. He therefore concludes that it was Omar. His colleague Kima Greggs, who is watching Lester via CCTV, knows that Bubbles was mistaken this time and mistook someone else for Marlo. Moreover, the equally reliable Johnny told Kima that he saw Marlo shooting Bob. Kima says to Detective McNulty, who is watching Lester via CCTV with her:

(4) Oh what a shame! Lester should believe that it was Marlo.

In fact, however, unbeknownst to Kima and McNulty, Johnny was also wrong and mistook Omar for Marlo. Omar indeed killed Bob.

The next day, Detective McNulty finds out about this and confronts Kima with the facts:

(5) Well, I guess you were wrong when you said that Lester should believe that it was Marlo.

Kima admits that she was wrong when she uttered (4).

Kima seems right to admit that (4) was wrong. However, she can only be right if the ‘should’ in (4) is objective. If it were collective, then she would rather respond to McNulty:

(6) No, given what we all knew at the time, Lester should have believed that it was Marlo!

To conclude, the fact that Kima is willing to retract (4), and does not refuse to, is evidence that the ‘should’ in (4) is objective, and not collective.

A third argument in favour of the objective ‘should’ is that ‘*S* should do *A*’, sometimes at least, means ‘It is best if *S* does *A*’.¹⁸ Take the example of the sentence ‘There should be world peace’. It arguably means that it would be best if there was world peace and expresses a political ideal. It is natural to describe

¹⁸See Sloman (1970). For more on the relation between ‘best’ and ‘should’, see section 2.4.

an epistemically ideal world as one in which we believe every true proposition and no false proposition (Gibbard 2005). This is connected to the point that the standard of correctness for belief is truth. In such an ideal world, we hold all the correct beliefs we could possibly hold. In other words, what is best, epistemically speaking, is to believe all and only true propositions. Thus, where ‘should’ means ‘best’, ‘One should believe P iff P ’ is true. Since this sentence can only be true if the ‘should’ occurring in it is objective, it follows that the doxastic ‘should’ has an objective sense. When speakers like Kim use the objective doxastic ‘should’, they describe what would be best in light of the epistemic ideal of truth. You do not necessarily have to believe that truth is the *only* epistemic ideal in order to be convinced by this third argument. You can believe that it’s just one epistemic ideal next to others such as knowledge, understanding, justification, or reliability, and that it is one of the epistemic ideals in light of which we evaluate doxastic attitudes.¹⁹

Fourth, as I have pointed out in section 3.2 above, it is widely agreed that the practical ‘should’ has an objective sense. Even though you are not aware that the hotel is on fire, you should, in some sense at least, jump out of the window. Now, in chapter 2, I explained that the doxastic and the practical ‘should’ are both instances of the deontic modal ‘should’. Whether this ‘should’ expresses epistemic or practical norms depends on whether the contextually salient ordering source reflects epistemic or practical norms. That practical and doxastic ‘should’ are the same except with respect to the kind of normativity they reflect makes it plausible that they also share the characteristic that they can have an objective sense.

Fifth, Schroeder (2015) suggests a subjective/objective distinction analogous to the one I am discussing here, but with respect to reasons. On his view, reasons are propositions. An objective reason is a true proposition. For example, the fact that Renaud is smiling is a reason for believing that Renaud is happy. A subjective reason is a proposition that the agent who has the reason has a certain access to, e.g., a proposition that is part of the agent’s evidence. For example, if I know that Renaud is smiling, I have a subjective reason to believe that he is happy. Now, the following principle is plausible:

¹⁹Alston (2005) proposes such a view.

Reasons-Should If there is conclusive reason for S to believe P , S should believe P .

Now, let's assume that P is true. Then, there is an objective reason to believe P for any person S : P itself. This might sound odd, but P clearly counts in favour of P . In fact, P seems to be a conclusive objective reason to believe P since there is no other fact, i.e., objective reason, that could outweigh it. Thus, we can infer from **Reasons-Should** that for any proposition P , if P is true, one should believe P . This conclusion can only be true if the 'should' occurring in it is objective. Since the derivation is sound, the conclusion is true and the doxastic 'should' must thus have an objective sense.

I have to concede that this argument is probably the weakest of the five. It rests on the assumption that there are objective reasons. If you need to be convinced that there is such a thing as an objective doxastic 'should', wheeling in objective reasons might feel close to begging the question.

3.4 Subjective Epistemic Norms and the Truth Norm

I will present the main thesis of this chapter in this section. Subjective epistemic norms (which I will introduce in subsection 3.4.1) and the truth norm (3.4.2) seem to be in conflict (3.4.3). I argue that this conflict is only apparent as formulations of the truth norm and those of subjective norms use different senses of the doxastic 'should', namely, the objective and subjective sense, respectively (section 3.4.4). As these senses result from the information-sensitivity of the doxastic 'should', one implication of this semantic feature for normative epistemology is that it resolves the apparent conflict between subjective and objective epistemic norms.

3.4.1 Subjective Epistemic Norms

There are a number of epistemic norms discussed in the literature that I would describe as *subjective*. What they all share is that they make what a subject should believe partly dependent on the subject's epistemic perspective. In section 3.2 above, I've already mentioned **Evidentialism** as an example. Another example is the knowledge norm, according to which one should, roughly, believe something only if one knows it. In his (2013: 14), Gibbons describes the knowledge norm as

“objective”. The reason is that due to the factivity of knowledge, the knowledge norm makes whether one should believe P dependent on an aspect of the external world: on whether P is true. Thus, for Gibbons (2013), a norm is subjective only if what one should believe is *entirely* dependent on the subject’s epistemic perspective, as for example with **Evidentialism**. For me, however, a norm is subjective if what one should believe is to *some extent* dependent on the subject’s epistemic perspective. Gibbons and I don’t have a substantive agreement; we’re just carving up logical space in different ways.

In this subsection, I want to show that due to its flexibility, **Subjective Doxastic Should**, and hence **Doxastic Should**, can accommodate subjective epistemic norms. This means, **Subjective Doxastic Should** can be precisified in a way that allows interpreting these subjective epistemic norms as using the subjective doxastic ‘should’ in their formulation. As a reminder, here is the account of the subjective doxastic ‘should’ I offered in section 3.2:

Subjective Doxastic Should S should subjectively adopt doxastic attitude D to P iff having D to P reflects accurately how much S ’s information-state supports P .

I already explained there how **Evidentialism** can be accommodated by **Subjective Doxastic Should**. The norm can be seen as an instance of the justification norm:

Justification S should have doxastic attitude D to P iff S is propositionally justified in adopting D to P .

Gibbons (2013: 5) explicitly proposes this norm. What it exactly amounts to depends on the notion of justification that the respective proponent has in mind. On Feldman’s evidentialist account of (propositional) justification **Justification** becomes **Evidentialism**.²⁰ Gibbons (2013: 39) propagates a more externalist account of justification according to which, roughly, one has propositional justification to believe something iff one’s available evidence supports it, where one’s available evidence is not just the evidence one already possesses, but

²⁰Hume also seems to support something along these lines, when he states: “A wise man [...] proportions his belief to the evidence” (Hume 1748/2007: sec. 10).

also evidence one can very easily collect, for example by looking around the room one is in. His version of **Justification** can be accommodated by **Subjective Doxastic Should**, too, if we take the subject's information-state to be their available evidence.

Another, very prominent account of justification is reliabilism.²¹ Unlike evidentialism, it is an account of doxastic justification. Roughly, one has propositional justification to have believe P iff one has sufficient reason to believe P . One has doxastic justification to believe P if one, furthermore, also actually believes P and does so for the right reasons (Kvanvig and Menzel 1990). A somewhat altered version of **Justification**, according to which one should believe a proposition iff one is doxastically justified in doing so, could look like this:

Reliabilism S should believe P iff believing P is the result of a reliable belief-forming process.

Kornblith (2002: 158) seems to defend such an account, when he proposes that epistemic norms require us to adopt cognitive systems that are truth-conducive. Precising **Subjective Doxastic Should** such that it makes **Reliabilism** true is a bit trickier, but possible. We can stipulate that the contextually provided information-state i is the proposition describing how reliable the process that produced S 's belief that P is. This might stretch the notion of 'information-state' a bit, but since it is a technical term, it need not refer to anything that looks like evidence. Furthermore, we stipulate that believing P reflects most accurately how much i supports P iff according to i the process that produced S 's belief that P is more than (let's say) 70% reliable. This would make sense since the probability of P on i , i.e., the probability of P given that the process that produced belief that P is more than 70% reliable, is more than 70%. Thus, P has an evidential probability of more than 70% on i , and this support relation seems to be most accurately reflected by believing P , rather than by suspending judgement on whether P or disbelieving P . (Of course, 70% is randomly chosen, and one can set the threshold required for licensing belief higher or lower.)

Admittedly, at this point one might wonder whether it is still accurate to speak of the information-sensitivity of the doxastic 'should' if the thing the doxastic

²¹See, for example, Goldman (1979).

‘should’ is sensitive to doesn’t have much to do with what we intuitively describe as information. This is not too worrisome for two reasons. First, it is far from clear that our intuitive understanding of technical terms like ‘information-state’ should bear on how we use them in our formal theories. Second, what is essential about the information-sensitivity of the doxastic ‘should’ is that it allows us to evaluate doxastic attitudes from different epistemic perspectives. One of these perspective is the perspective of the relevant subject, and this is the perspective we’re taking when we employ the subjective doxastic ‘should’. Now, an obvious candidate for characterizing a subject’s epistemic perspective is their (available) evidence. However, the reliability of the process that brought about a certain belief is also an aspect of the person’s epistemic perspective. Imagine that we both believe that Germany won the Football World Cup in 2014, but that you do so because you read it in the newspaper and I because of wishful thinking. The different ways in which our beliefs were formed puts us in different epistemic positions with respect to the proposition that Germany won the Football World Cup in 2014. To conclude, the proposition that describes how reliable the process is that formed a subject’s belief that P represents the subject’s epistemic perspective on P and can therefore be called the subject’s information-state (with respect to P).

Another, quite popular subjective norm is the knowledge norm. If we were to simply model it on **Justification**, it would read as follows:²²

Knowledge-Simple S should believe P iff S knows that P .

Knowledge-Simple has the problem that since knowledge entails belief, it follows from it that if one should believe P , one believes P . This is obviously wrong because we’re not epistemically perfect and can fail to believe what we should (Whitcomb 2014: 93).²³ To avoid this, one could weaken **Knowledge-Simple** as follows:

²²The following remarks on the knowledge norm owe much to the excellent discussion in Whitcomb (2014).

²³Notice that **Reliabilism** faces the same problem since it entails that we should believe P if we are doxastically justified to believe P , from which follows that we should believe P only if we believe P . Defenders of **Reliabilism** could react to this by modifying their position such that one should believe P iff if one came to believe P , one’s belief that P would be the product of a reliable belief-forming process. This position avoids the problem of epistemic perfection for similar reasons as **Knowledge-Counterfactual**, which I discuss below.

Knowledge-Material S should believe P only if S knows that P .²⁴

Knowledge-Material replaces the biconditional in **Knowledge-Simple** with a material implication. As a result, **Knowledge-Material** does not entail that we're epistemically perfect. However, in turn this norm never tells us when we should believe P or should not believe P . It only tells us when *it is not the case* that one should believe P : when we don't know that P . Notice, \langle it is not the case that S should believe P \rangle is not the same as $\langle S$ should not believe P \rangle .²⁵ The former can be true and the latter false if, for example, S should either believe P or suspend judgement on P , but both options are permissible.

The most common version of the knowledge norm is even weaker than **Knowledge-Material**:

Knowledge-Wide S should: believe P only if S knows that P .²⁶

To understand the distinction between the two, it helps to look at their logical form, which is respectively **LF-Knowledge-Material** and **LF-Knowledge-Wide**:

LF-Knowledge-Material Should: (S believes P) only if S knows P .

LF-Knowledge-Wide Should: (S believes P only if S knows P).

'Should' has narrow scope in **Knowledge-Material**, while it has wide scope in **Knowledge-Wide**. In the former, it binds the expression ' S believes P ', in the latter the whole expression ' S believes P only if S knows P '. As Whitcomb (2014: 94) points out, such a wide-scope formulation of the knowledge norm is even more uninformative about what we should believe than **Knowledge-Material**. It only tells us that we should not be in a state where we believe P , but do not know P . This norm can be satisfied in two ways: either by not knowing and not believing P or by knowing P and believing P .²⁷ **Knowledge-Wide** is thus not only incapable of telling us whether we should or should not believe a proposition,

²⁴This norm is discussed, but not endorsed, by Littlejohn (2012: 233).

²⁵In terms of modal logic, this is the difference between $\neg\Box$ and $\Box\neg$.

²⁶See Williamson (2000: 11), Sutton (2007: 44), and Bird (2007: 95).

²⁷The third option would be knowing and not believing P , but since knowledge entails belief, this is a logical impossibility.

but also of telling us whether it is not the case that one should (or not) believe a proposition.²⁸

To sum up, **Knowledge-Material** and **Knowledge-Wide** are strangely ‘hands-off’ for a norm of belief. Another way to circumvent the difficulty with **Knowledge-Simple** is to go counterfactual:

Knowledge-Counterfactual S should believe P iff if S came to believe P , S would know P .²⁹

This does not entail that if one should believe P , one believes P , since according to it one should believe P if one does not believe P , but would know it if one came to believe it. Furthermore, it clearly tells us under which conditions we should (not) believe a proposition. Whitcomb (2014) points out that **Knowledge-Counterfactual** faces some counterexamples due to the conditional fallacy (Shope 1978) and considers different possible moves to fix this problem. I don’t want to go into this here; I merely wanted to lay out that there is a version of the knowledge norm that neither entails that we’re epistemically perfect nor is too hands-off. In the following, if not otherwise indicated, I will orientate myself on **Knowledge-Counterfactual** when discussing the knowledge norm.

Subjective Doxastic Should can accommodate **Knowledge-Counterfactual** as follows. First, we need to assume that information-states are not sets of worlds, but sets of propositions. Furthermore, a subject S ’s information-state is the set of the propositions that S would come to know if S believed them. Let’s call this set of propositions S ’s *counterfactual knowledge*. That S ’s counterfactual knowledge is S ’s information-state is somewhat intuitive. On the Williamsonian assumption that $E = K$, i could then be described as S ’s *available* evidence, since it is the knowledge, i.e., the evidence, that S either already possesses or can gather. Finally, we stipulate that a belief that P accurately reflects how much i supports P iff i fully supports P . ‘Full support’ is a technical notion such that an information-state i fully supports P iff P is a member of i . It makes intuitive sense to see this as the strongest possible evidential support relation

²⁸Whitcomb’s argument is inspired by Broome’s (1999) argument that principles of rationality are not normative since they are only wide-scope obligations.

²⁹This is akin to another natural formulation of the knowledge norm according to which one should believe P iff one is in a position to know P . Smithies (2012) puts this norm forward.

between a proposition and an information-state; the proposition is not just very likely on the information-state, it is actually *contained* by it. If we stipulate that belief accurately reflects only such full support and that i in **Subjective Doxastic Should** is S 's counterfactual knowledge, then S should subjectively believe P iff if S came to believe P , S would know P .

The justification and knowledge norm are not the only subjective norms one can hold. For one, one can defend **Evidentialism** or **Reliabilism** without accepting evidentialist or reliabilist accounts of justification. In this case, one would be non-committal with respect to **Justification**, but such epistemic norms would of course still be subjective. Furthermore, the principle of conditionalization championed by Bayesian epistemology is equally subjective as it tells you how to revise your credences in the light of new evidence you have gathered. This can easily be accommodated by **Subjective Doxastic Should** if we precisify it such that (i) i is the subject's total evidence and that (ii) a credence of x in P accurately reflects i 's support for P iff x is the conditional probability of P on i .

To sum up, it holds for all of the subjective epistemic norms discussed here, that **Subjective Doxastic Should** is flexible enough to be precisified in such a way that it makes them true in the subjective sense of the doxastic 'should'. Before I move on to the objective norm, it is worth noting that the different subjective norms are also in conflict with each other. For example, one can justifiably believe, but not be in a position to know P . In this case one should believe P according to the justification norm, but not the knowledge norm. One can approach the conflict between the subjective norms by claiming that only one of them is correct and that the precisification of **Subjective Doxastic Should** that accommodates this one norm gives us the one correct meaning of the subjective doxastic 'should'. Alternatively, one could argue that all (or at least some) of the subjective norms are true in *some* sense of the subjective 'should'. On this picture, the information-state that the subjective 'should' is relativized to is, for example, on some occasions the subject's evidence and on other occasions it is the subject's counterfactual knowledge. In effect, this would mean that the doxastic 'should' has several subjective senses. My goal here is not to resolve the conflict between subjective epistemic norms and to decide whether there is one subjective

sense of the doxastic ‘should’ or in fact several.³⁰ I am only concerned with the (apparent) conflict between subjective and objective norms. When I speak of *the* subjective doxastic ‘should’ in the following, this is therefore *not* supposed to mean that I side with the view that there is only one subjective sense, but is just shorthand.

3.4.2 The Truth Norm

Let’s turn to the truth norm. In this subsection, I will quickly present the different versions of the truth norm that can be found in the literature, what motivates some of them, and how they can be accommodated by **Objective Doxastic Should**. The default formulation of the truth norm is the one I used above:

Truth For any proposition P and subject S , S should believe P iff P is true.³¹

Gibbard (2005), Shah and Velleman (2005), and Engel (2007) are among the proponents of **Truth**. As I have already shown in section 3.2, it can be accommodated by

Objective Doxastic Should For any proposition P and subject S , S should objectively believe P iff P is true.

However, **Truth** is not the only formulation of the truth norm on offer. Wedgwood (2002: 273) proposes the following:

Truth-Considered For any subject S and any proposition P that S considers, S should believe P iff P is true.

Whereas **Truth** implicitly quantifies over all propositions, **Truth-Considered** only quantifies over those propositions that the subject considers. This is to avoid the objection to **Truth** that it would make more sense to suspend judgement on trivial propositions, such as those concerning the number of grains of sand on

³⁰However, in chapter 4 I propose that we should distinguish between two senses, one that is relativized to the subject’s first-order evidence and one that is relativized to their higher-order evidence.

³¹In the following, I will make the quantification over propositions and subjects explicit again, since they are important for our discussion.

a beach, rather than to believe them since believing them would be a waste of one’s valuable cognitive resources. Since on **Truth-Considered** the principle only applies to propositions one considers anyway and thus already spent cognitive resources on, **Truth-Considered** is not subject to this objection.

The ‘should’ in **Truth-Considered** cannot be objective given **Objective Doxastic Should** since according to the latter, ‘*S* should believe *P*’ is true even if *P* is not considered by *S* as long as *P* is true. A possible response to this would be to modify **Objective Doxastic Should** as follows:

Objective Doxastic Should-Considered For any subject *S* and any proposition *P*, *S* should objectively believe *P* iff *P* is true and *S* considers *P*.

This refined account of the meaning of ‘*S* should believe *P*’ where the ‘should’ in it is the objective doxastic ‘should’ quantifies over all subjects and propositions, unlike **Truth-Considered**. This is necessary because otherwise it wouldn’t give us an account of the meaning of all sentences of the form ‘*S* should (objectively) believe *P*’. Despite this difference, **Objective Doxastic Should-Considered** entails **Truth-Considered** and hence makes it true. **Objective Doxastic Should-Considered** can be integrated into **Doxastic Should** by stipulating that where the doxastic ‘should’ takes on an objective sense the contextually provided information-state is the intersection of all those true propositions which are such that the subject *S* considers them. I take it that if one considers a proposition *P* one thereby automatically also considers its negation. Thus, if someone wonders ‘Is *P* true?’ and *P* is false, then $\neg P$ is among the true propositions considered by this person. This information-state will have a strong support relation with the true propositions *S* considers since it entails them, and will neither support nor speak against those propositions *S* does not consider. Thus, it will license believing *P* iff *P* is true and considered by *S*.

Another concern with **Truth** is raised by Boghossian (2003). He argues that it cannot hold that for every true proposition, we should believe it. This would violate the principle that ‘ought’ implies ‘can’ since it is impossible for someone to believe everything that is true—this would be a cognitive overload. He therefore retreats to the following, weaker formulation:

Truth-Necessity For any proposition P and subject S , S should believe P only if P is true.³²

A problem with this account, just like with **Knowledge-Wide**, is that it does not tell us when one should believe a proposition (Bykvist and Hattiangadi 2007: 280). If P is true, nothing follows about what S should believe from **Truth-Necessity** since it only gives us a necessary condition. If P is false, it only follows that it is not the case that S should believe P , but not that S should not believe P . For this reason, I will mostly ignore it in the following and hence don't need to discuss whether it can be accommodated by **Objective Doxastic Should**.

Notice that **Truth-Considered** is not subject to the point that 'ought' implies 'can' (Bykvist and Hattiangadi 2007: *ibid*). It is within our cognitive capacity to believe all true propositions *that we consider*. Wedgwood's (2002) account of **Truth** is at the same time logically strong enough to make predictions about what one should believe. I will discuss whether **Truth** or **Truth-Considered** is to be preferred in section 3.5. For now, what matters is that they can be accommodated by either **Objective Doxastic Should** or **Objective Doxastic Should-Considered**, which can both be integrated into **Doxastic Should**.

3.4.3 The Conflict

Subjective epistemic norms and the truth norm seem to conflict with each other. More precisely, they appear to be inconsistent. **Truth** and **Justification** can clash in two ways. First, on almost any conception of justification, there are justified false beliefs.³³ Thus, there can be a situation where the following is true:

(7) S is propositionally justified to believe P and P is false.

From (7) and **Justification**, we can infer:

(8) S should believe P .

From (7) and **Truth**, we can infer:

(9) It is not the case that S should believe P .

³²Shah (2003: 469f.) also chooses this formulation of the truth norm.

³³Littlejohn (2012) opposes this traditional view.

Notice that (9) can also be inferred from (7) and **Truth-Considered**, given that *S* considers *P*. (8) and (9) are inconsistent. Thus, **Justification** and **Truth** (and its variants) clash if (7) holds.

Second, there are cases where (10) holds:

(10) *S* is not propositionally justified to believe *P* and *P* is true.

Justification and (10) entail:

(11) It is not the case that *S* should believe *P*.

Truth and (10) entails:

(12) *S* should believe *P*.

(11) and (12) are inconsistent, so **Justification** and **Truth** clash where (10) holds. The same is true for **Justification** and **Truth-Considered** if *S* considers *P*. It will be easy to see that these two kinds of conflict can also occur between the truth norm and all the other subjective epistemic norms I discussed above, except for the knowledge norm. For example, it is possible that one's evidence supports believing *P* and *P* is false or that one's evidence supports not believing *P* and *P* is true. In those cases, **Evidentialism** and **Truth** clash.

Variations of the knowledge norm and the truth norm can come in conflict, too. We can ignore **Knowledge-Narrow** since it is extremely implausible and probably no one holds it anyway. However, **Knowledge-Counterfactual** and **Truth** clash when (13) holds:

(13) If *S* came to believe *P*, *S* wouldn't know *P*, and *P* is true.

On an analysis of knowledge as un-Gettierized justified true belief, (13) holds when *S* is not propositionally justified to believe *P* or *S*'s belief that *P* would be Gettierized.

(13) and **Truth** entail:

(14) *S* should believe *P*.

(13) and **Knowledge-Counterfactual** entail

(15) It is not the case that S should believe P .

Thus, when (13) holds, **Knowledge-Counterfactual** and **Truth**, as well as **Truth-Considered** where S considers P , clash.

What about other instances of the truth norm and the knowledge norm? Since **Knowledge-Wide** does not entail anything about whether it is (not) the case that one should (not) believe a proposition, it cannot be in conflict with **Truth** or any variation of it. **Truth-Necessity** only tells us something about whether a subject S should believe a proposition P when P is false, in which case it says that it is not the case that one should believe P . However, the knowledge norm tells us in none of its variations that one should believe P when P is false, so **Truth-Necessity** and the knowledge norm can never conflict.

To sum up, subjective epistemic norms and the truth norm seem to be inconsistent, except if we choose **Knowledge-Wide** as our subjective norm or we choose any instance of the knowledge norm as our subjective norm and **Truth-Necessity** as our preferred formulation of the truth norm. As I have mentioned, both **Knowledge-Wide** and **Truth-Necessity** have been criticized for being too weak since they don't tell us what we should believe, which one would guess is the precise job of an epistemic norm. Accordingly, I will put **Knowledge-Wide** and **Truth-Necessity** aside in the remainder of this chapter. Assuming that we want stronger epistemic norms, we seem to face a puzzle. As Gibbons (2013: 6) points out, subjective norms and the truth norm both seem to have something going for them, but if they are inconsistent, we can't accept both.

3.4.4 Resolving the Conflict

The distinction between the subjective and the objective doxastic 'should' provides a straightforward solution to the conflict between the two kinds of norms. Any of the apparently inconsistent pairs of sentences that we derived from variants of the truth norm and the subjective norms are not actually inconsistent. For example, (8) and (9) explicitly mean:

(8*) S should *subjectively* believe P .

(9*) It is not the case that S should *objectively* believe P .

(8*) and (9*) are not inconsistent, just as the sentences ‘This is a bank’ and ‘It is not the case that this is a bank’ are not inconsistent if ‘bank’ in the first sentence refers to a river bank and ‘bank’ in the second to a financial institution.

3.5 Four Objections

While probably everybody who believes that there are epistemic norms would agree that at least *some* subjective norm holds, the truth norm is significantly less popular. An alternative solution to the conflict between subjective norms and the truth norm is therefore to simply give up the latter. Furthermore, claiming that the truth norm is not correct is claiming that there is no objective doxastic ‘should’ as the correctness of the truth norm entails that there is an objective doxastic ‘should’. In the following, I will first discuss three objections to the truth norm, which if successful, would accordingly undermine my solution. Subsequently, I will reply to a fourth objection by Gibbons (2013) that directly attacks the idea that there is a subjective and an objective sense of the doxastic ‘should’.

First, as I already mentioned in section 3.4, one of the major worries with the truth norm is that following it will either lead to cognitive overload or at least to an inefficient use of our cognitive capacities.³⁴

A first reply to this worry is that the truth norm describes an epistemic ideal. As mentioned in section 3.3, the objective doxastic ‘should’ can be seen as describing what is epistemically best. When we think about ideal situations, we abstract away from real-world constraints. For example, given the current political situation or maybe even given the human condition, it might be impossible to achieve world peace. ‘There should be world peace’ could nonetheless be true since it would be best if there was world peace. Analogously, when describing what would be epistemically best we presumably ignore constraints such as our cognitive capacities, too (see Engel 2013: 209). Thus, ‘One should believe all and only all true propositions’ can be true even if fulfilling this is psychologically impossible. As Gibbard (2005) nicely points out, this abstraction away from real-life constraints also applies to subjective norms like the normative constraints of probabilism. The latter require that we get rid of all probabilistic incoherencies in our credences, like assigning a proposition and its negation a total credence

³⁴See Wedgwood (2002), Boghossian (2003), and Bykvist and Hattiangadi (2007).

of less or more than 1. However, Gibbard is certainly right when he writes that “[o]n many matters [...] ironing out incoherencies would be more trouble than it’s worth” (ibid.: 339-40).

A second, alternative reply to this worry about cognitive overload and waste is Wedgwood’s (2002) suggestion to restrict the truth norm to propositions one has considered. As explained above, I can accommodate this move by proposing **Objective Doxastic Should-Considered** as the definition of the objective doxastic ‘should’.

I am not sure which of these two general strategies are to be preferred. I won’t settle this issue here and will remain neutral with respect to the question whether the objective doxastic ‘should’ is better characterized by **Objective Doxastic Should** or **Objective Doxastic Should-Considered**.

Second, Bykvist and Hattiangadi (2007: 281) argue that the truth norm faces difficulties because of so-called *blindspot propositions* like the following one:

(16) It is raining and nobody believes it is raining.

If (16) is true, you should believe it according to **Truth**, but then (16) becomes false because now someone believes that it’s raining: you. The problem is not that you cannot believe (16) if you should believe it. Even though it is irrational, it seems psychologically possible to form the belief that it is raining and that nobody believes that it is raining. It is just that as soon as you try to satisfy the truth norm and form the belief that (16), it’s no longer the case that you should believe it as it has become false. **Truth** hence violates the plausible principle that if one should do *A*, one can do *A* while it is the case that one should do *A*. Notice that the same problem applies to **Truth-Considered** since we can easily imagine that you’re considering (16).

As Bykvist and Hattiangadi (2007: 282) themselves point out, one way to address this problem is to restrict **Truth** to propositions unlike (16) that can be believed truly, so that **Truth** doesn’t entail that we should believe (16):

Truth-Believable For any subject *S* and any proposition *P* that is truly believable for *S*, *S* should believe *P* iff *P* is true.

In this case, we could modify **Objective Doxastic Should** as follows in order

for it to fit **Truth-Believable**:

Objective Doxastic Should-Believable For any subject S and proposition P , S should objectively believe P iff P is true and P is truly believable for S .

Like **Objective Doxastic Should-Considered**, **Objective Doxastic Should-Believable** quantifies over all subjects and propositions, so that it can provide the truth-conditions for every sentence of the form ‘ S should (objectively) believe P ’. It entails **Truth-Believable** and thus makes it true.

Modifying **Objective Doxastic Should** in this manner is not *ad hoc*. If we assume that the objective doxastic ‘should’ describes the epistemic ideal of a world where we hold every correct belief we can possibly hold, it makes sense to exclude propositions from the range of the truth norm that are such that it’s impossible to believe them correctly. Notice that ‘impossible’ here means *logically* impossible (ibid.: 281). I’ve argued above that it need not be a problem if the truth norm describes an ideal that is *humanly* impossible to achieve as it would lead to a cognitive overload. Ideals don’t have to be achievable. This is not incoherent with the claim that the truth norm shouldn’t apply to propositions that aren’t truly believable. While ideals might ignore human capacities, they can’t ignore what’s logically possible.

Doxastic Should can accommodate **Objective Doxastic Should-Believable**. However, this is a bit more complicated. We need to assume that the objective doxastic ‘should’ is relativized to an information-state i representing true propositions that can be believed truly by S . Let’s call this the *truly believable realistic* information-state, or i_{tbr} . One important point is that in order for **Doxastic Should** to accommodate **Objective Doxastic Should-Believable**, i_{tbr} must not be modelled as the intersection of all these propositions, where the propositions are again modelled as sets of worlds. Otherwise, we end up with i_{tbr} being a set of worlds such that in all of them (16) holds true, given that the conjuncts of (16) are true and truly believable. Because a set that is an intersection of these two conjuncts and other propositions will only contain worlds in which both conjuncts are true, such a set will only contain worlds in which (16) is true. If i_{tbr} was modelled in this fashion, it would provide strong evidential support for (16) and hence license believing (16). In this case, **Doxastic Should**

would predict that ‘You should believe (16)’ is true on an objective reading of ‘should’.

To avoid this, we can model i_{tbr} as a set of propositions, and not intersections of them. Furthermore, we need to stipulate that believing P accurately reflects the support relation between i_{tbr} and P iff i_{tbr} fully supports P . I have defined the notion of ‘full support’ in subsection 3.4.1 such that an information-state i fully supports a proposition P iff i is a member of P . This way, one should believe P in light of i_{tbr} , i.e., objectively, iff P is true and can be truly believed. The same solution applies *mutatis mutandis* to **Truth-Considered** and **Objective Doxastic Should-Considered**.³⁵

Bykvist and Hattiangadi (2007: 282) object that formulations of the truth norm along the line of **Truth-Believable** are too weak since they do not tell us what to do if we’re faced with blindspot propositions. Wedgwood (2013: 128), who defends **Truth-Considered**, agrees with this. For him, the truth norm is the foundational epistemic norm (Wedgwood 2002) that all other epistemic norms derive from. Hence, the truth norm needs to cover every proposition. This is not really a problem for my view. All I claim is that the truth norm holds since the ‘should’ in its formulation is the objective doxastic ‘should’. I do not make any claims about the position of the truth norm in the hierarchy of epistemic norms. I do not need it to cover every proposition.

Admittedly though, ‘One should not believe (16)’ has a true ring to it. However,

³⁵One might worry whether it’s still possible for **Doxastic Should** to account for both the subjective *and* the objective doxastic ‘should’ if the notion of ‘accurately reflecting a support relationship’ is defined in such a peculiar manner. If believing P accurately reflects the support relationship between an information-state i and P only if P is a member of i , how can **Doxastic Should** explain that ‘ S should (subjectively) believe P ’ is true where S ’s evidence abductively supports P , but P is not among the propositions constituting S ’s evidence? Given our definition of ‘accurately reflecting’, S ’s evidence would not license believing that P .

This problem reveals a limitation of **Doxastic Should** that arises from its simplified nature. As I made clear in section 2.7, **Doxastic Should** is a simplified version of those semantic theories of the deontic ‘should’ that make use of information-sensitive ordering sources. While **Doxastic Should** struggles with the just described problem, these more sophisticated models can resolve it. We only need to design our information-sensitive doxastic ordering source g in the following manner: where g is relativized to the truly believable realistic information-state i_{tbr} , it ranks a world at which the subject holds a belief that P best iff i_{tbr} fully supports P , i.e., iff P is a true proposition that is truly believable. However, where g is relativized to a subject’s body of evidence E , g ranks in a more ‘relaxed manner’; it ranks a world where the subject holds a belief that P best also if E supports P less than fully, for example where E makes P only probable.

To conclude, **Doxastic Should** runs into difficulties accommodating both **Subjective Doxastic Should** and **Objective Doxastic Should-Believable**, but this is only so because it’s a simplified semantic account. This is not a problem for Doxastic Contextualism in itself.

I can make it turn out true by the following further development of the assumption that the doxastic ‘should’ is objective where it is relativized to the truly believable realistic information-state, i_{tbr} . I can stipulate that i_{tbr} does not fully support a proposition P , suspending judgement accurately reflects the support relation between i_{tbr} and P . We then end up with the view that for all true propositions P , if they are truly believable, you should objectively believe them, and if they aren’t truly believable, you should objectively suspend judgement about them. This seems plausible. If P is true, disbelieving it seems wrong in light of the ideal of truth, but if it can’t be truly believed, believing it seems wrong, too. Suspending judgement appears to be the best way out.

To conclude, we can defend the truth norm against the worry about blindspot propositions by restricting it to **Truth-Believable** and **Doxastic Should** can be specified such that it makes **Truth-Believable** true.

A third major objection to the truth norm is that it cannot be guiding and thus cannot really be a norm. Glüer and Wikforss (2009, 2010, 2015) argue that in order to follow the truth norm, we need to figure out first whether the proposition P we are considering is true or not. In other words, in order to follow it, we first need to form a belief about whether P holds, which is “the very belief the formation of which [the truth norm] was supposed to guide” (Glüer and Wikforss 2010: 752). The truth norm’s guidance therefore “necessarily comes too late” (ibid.: 751).³⁶ Notice that the argument is *not* that the truth norm cannot be guiding since one is not always in a position to tell whether the condition for its application is satisfied or not. This is a popular internalist-spirited criticism of any norm whose satisfaction conditions are such that we are not always in a position to know whether they apply. However, as Boghossian (2003) remarks, among such norms is not only the truth norm, but many other norms, such as ‘Buy low, sell high’. Furthermore, Williamson (2000: 192) has argued that since no interesting condition is luminous, i.e., such that we are always in a position to know whether it obtains or not, actually no norm can be such that we are always in a position to know what we need to do in order to conform with it.

This does not undermine the defence of the truth norm that I’ve been offering here. I am arguing that there is a reading of the truth norm on which it is true—i.e., when the doxastic ‘should’ in its formulation is objective. I’m not

³⁶Sosa (2001: 54) makes the same argument.

claiming that this ‘should’ is guiding. As McHugh (2012) has argued, we do not need to interpret the truth norm as being prescriptive, i.e., as being “apt to guide an agent in doxastic deliberation” (ibid.: 18). He proposes to understand the truth norm instead as evaluative and therefore best formulated in terms of ‘good’, not ‘should’. More precisely, he formulates the truth norm as follows (ibid.: 19):

Truth-Good For any subject S and proposition P , if S believes P , then that belief is good if P is true and bad if P is false.

My proposal differs in a number of ways. First, I think the truth norm can be formulated in terms of ‘should’. This is arguably only a difference in terminology since on my proposal the objective doxastic ‘should’ means something similar to ‘good’, namely ‘best’, and is not prescriptive. Second, similar to Wedgwood (2002), McHugh (2012: 21) claims that the truth norm needs to explain other epistemic norms. As I have pointed out, I am not committed to this view. Finally, on McHugh’s account the truth norm only evaluates propositions a subject believes, whereas on my picture it applies to all propositions that they can truly believe (or at least all propositions that they consider and can truly believe). McHugh’s main reason for rejecting more inclusive formulations like **Truth-Believable** seems to be that he believes that even if the truth norm does not guide our belief-forming processes directly, it nonetheless plays a motivating role in these processes. For example, he suggests that our acceptance of the truth norm should explain why we avoid believing propositions for which we lack evidence (ibid.: 19). The trouble with **Truth-Believable** is that it is too complicated to play such a role (ibid.: 15). Again, I do not ascribe to any particular view on how the truth norm is connected to other epistemic norms or our cognitive processes. On the picture I suggested in section 3.3, we use the objective doxastic ‘should’ to describe which beliefs others persons should adopt in light of the epistemic ideal of truth.³⁷ While **Truth-Believable** might be too complicated to play a role in the deep-seated psychological processes of belief-formation, it does not seem to be too complex to be employed for the more conscious practice of evaluating other people. If it were, then many analyses of knowledge that go beyond the simple JTB-analysis

³⁷Sometimes this third ‘person’ and the evaluated subject are the same, but at different points in time.

would disqualify just on basis of being too complicated to explain our practice of ascribing knowledge to each other. However, despite all these differences, I agree with McHugh's (2012) central point: the truth norm needs not to be read as prescriptive. Therefore, the point that the truth norm cannot guide does not show that the norm can't be true.

A fourth complaint with resolving the conflict between the two kinds of norms by arguing that they use different senses of the doxastic 'should' comes from Gibbons (2013). He acknowledges that 'should' might be ambiguous or context-sensitive in the way I have described it, but then adds that he does not think that this can "really solve the problem" (ibid.: 19). For Gibbons, the question remains which of the two expresses a "genuine requirement" (ibid.: 21) or the "regular ought" (ibid.: 65). This is in and of itself not much of an argument. It is as if you pointed out that my question 'Are banks close to water?' is ambiguous, and I replied by saying 'OK. But are *genuine* banks close to water?'

Gibbons has more to say in favour of his point. He remarks that practical deliberation has the aim of making up one's mind, of coming to a practical conclusion by forming an intention (ibid.: 62). In practical deliberation, we are guided by the question 'what should I do?', and this question must be unambiguous, otherwise it couldn't lead us to a specific practical conclusion. Applying this to doxastic deliberation, Gibbons argues that the core issue is whether the subjective or the objective doxastic 'should' figures in the question 'What should I believe?' that leads us in doxastic deliberation. Only one of the senses of the doxastic 'should' can be the one that guides deliberation, otherwise doxastic deliberation would be paralysed. Gibbons is not the only one to hold this position. In his review of Gibbons (2013), Lord (2014) agrees that the ambiguity theory can for this reason not solve the conflict between subjective norms and the truth norm.

In other words, Gibbons claims that there can only be one norm that is guiding us and thus either the truth norm or a subjective norm is correct, but not both. As the reader will have guessed, my reply to this is that I happily grant that it is one of the subjective norms that is guiding, but that the truth norm is nonetheless correct as its purpose is not to guide, but to describe an epistemic ideal. One might have quarrels with calling the truth norm a *norm* or the objective doxastic 'should' *normative* if they are not involved in guiding doxastic deliberation. However, as McHugh (2012: 26, n. 12) notes, this terminological difference is irrelevant as

long as it is made clear that we're using the term 'norm' as an umbrella term, covering both guiding (or prescriptive) and evaluative norms.

3.6 Conclusion

The distinction between the subjective and the objective 'should' is familiar in ethics and has recently received more attention in epistemology. It can be seen as an instance of the information-sensitivity of the doxastic 'should' and is easy to integrate into the model of the latter that I have introduced in chapter 2. While the claim that there is an objective sense of the doxastic 'should' might be met with criticism, I have given five arguments in support of it, in particular instances of our use of it in natural language, where we employ it to assess beliefs in light of their standard of correctness. The pay-off of this distinction between the subjective and the objective doxastic 'should' for epistemology is that it solves the apparent conflict between the truth norm and subjective epistemic norms.

Chapter 4

Higher-Order Evidence

4.1 Introduction

In recent years, there has been a considerable amount of work in epistemology on higher-order evidence (HOE) and how an agent should react to such evidence. This question is particularly pressing in cases where the gathered HOE is misleading. For example, HOE might indicate that the proper functioning of an agent's cognitive faculties was interfered with when forming their belief that P on the basis of their first-order evidence (FOE), even though they actually inferred correctly from their FOE. One kind of HOE that can put us in such a scenario is peer disagreement,¹ which has received even more attention, compared to HOE in general, over the last 10 years or so.

Weatherson (ms) and Lasonen-Aarnio (2014) defend the view that one should stick with one's FOE in light of misleading HOE in general, and Van Inwagen (1996) and Kelly (2005) defend this view only for the particular case of peer disagreement. Borrowing a term used in the peer disagreement debate, let's call views according to which one should follow one's FOE rather than one's HOE *Steadfastness*.²

On the other side, Christensen (2010), Feldman (2009), and Schoenfield (ms) argue that our HOE puts rational pressure on us to sway away from our FOE.

¹See Feldman (2009), Christensen (2010), Kelly (2010), Lasonen-Aarnio (2014), and Schoenfield (2015).

²One could potentially add Lackey (2010) to this list. She defends the weaker claim that at least in some cases one should maintain one's FOE-supported belief, namely those where this evidence gives one a high degree of justification.

Elga (2007) makes this claim with respect to peer disagreement. Let's call this family of views *Conciliationism*.

When there is disagreement about which of two apparently conflicting positions is correct, there is a third position we can take. We can claim that actually both of the positions are correct: there is an appearance of conflict without a genuine conflict. I defend this option with respect to the HOE debate between Steadfastness and Conciliationism. In particular, I will argue in this chapter that an agent should both follow her FOE *and* follow her HOE, but just in different senses of the term 'should'. The choice between Conciliationism and Steadfastness is a false dichotomy.

Following up on the distinction between the subjective and the objective doxastic 'should' in the previous chapter, I suggest that the subjective doxastic 'should' can be divided further. We don't only distinguish in English between what an agent should believe in light of their evidence and in light of the relevant facts, but we can also make a more fine-grained distinction concerning the subjective 'should': between what an agent should believe in light of their FOE and in light of their HOE. In the following, I refer to these respective uses of 'should' as 'should_{FOE}' and 'should_{HOE}'.

In line with this distinction, my thesis in this chapter is that there is no such thing as what an agent faced with misleading HOE should believe *simpliciter*. Rather, we have to distinguish between what one should_{FOE} and what one should_{HOE} believe. Consequently, we do not have to be either Conciliationists *or* proponents of Steadfastness—we can have our cake and eat it, too. I call this position the *Evidence-Ambiguity Theory*.

Section 4.2 introduces the notion of higher-order evidence and section 4.3 some arguments in support of Steadfastness and Conciliationism. In section 4.4, I make the distinction between what one should_{FOE} and should_{HOE} believe more precise by integrating it into the information-sensitive contextualist semantics for the doxastic 'should' I have presented in chapter 2. Accordingly, I take the HOE debate to be a further case where Doxastic Contextualism helps to solve a tricky problem in epistemology.

Section 4.5 describes a potential worry with the Evidence-Ambiguity Theory: agents who are deliberating theoretically ask themselves the question of what they should believe. But it does not seem as though the Evidence-Ambiguity

Theory tells us what the answer to their question is since the question appears to be ambiguous on the view. In section 4.6, I deal with this objection by explaining how, on my account, the context of a deliberating agent provides a particular information-state and thus determines a particular meaning for the deliberating agent's question. Finally, in section 4.7, I distinguish the Evidence-Ambiguity Theory from the similar looking, but distinct position that misleading HOE puts us in a rational dilemma.

4.2 What Is Higher-Order Evidence?

Higher-order evidence (HOE) is, roughly, evidence about evidence. In particular, n^{th} -order evidence is evidence about $(n-1)^{\text{th}}$ -order evidence. In this chapter, we will focus on "second-order evidence" (Feldman 2005: 100), i.e., HOE about first-order evidence. First-order evidence (FOE) bears directly on the proposition the agent is wondering about (Kelly 2014: sec. 2). I will call this the 'target proposition'.

Evidence can be about FOE in different ways. It can be about what FOE other people possess or is available to you, but which you yourself don't possess (Feldman 2009: 304). This is not what I take to be HOE in this chapter. Rather, an agent's HOE, as understood here, bears on which proposition is supported by this agent's FOE. It can do this in different ways. First, it can be directly about the evidential relation that holds between the agent's FOE and a proposition P , for example when someone tells the agent that their FOE supports that Obama is the President of the US. Second, it can concern the reliability of the agent's cognitive faculties (Feldman 2005: 96; Kelly 2010: 139; Kelly 2014: sec. 2). Evidence of the latter kind in combination with what judgement one has formed on the basis of one's FOE is evidence to the effect of what one's FOE supports. Third, it can concern what FOE one has (Worsnip 2015b: 17).

Misleading HOE is similar to the phenomenon of ordinary undercutting defeat, however many, nonetheless, consider the two to be distinct (Feldman 2005: 113; Christensen 2010: 195; Lasonen-Aarnio 2014: 317). I agree that there is a distinction here. What I say in this chapter about how people should react to misleading HOE does therefore not extend to undercutting defeat. A classic example of ordinary undercutting defeat comes from Pollock and Cruz (1999: 86).

Suppose you look at an object that seems red to you. You are then justified to believe it is red. Then you learn that the room is bathed in red light and you know that under such circumstances objects that are actually white can seem red. You thereby lose justification for believing that the object is red. Your evidence that the room is bathed in red light is an undercutting defeater for your (propositional) justification for believing that the object is red.

It might appear as though this is simply an instance of HOE as I have characterized it. After all, it seems to suggest that your FOE does not support that the object is red. But I think this is false. True, your original FOE *in combination* with the evidence that the room is bathed in red lightning does not support that the object is red (or at least less so). But the undercutting defeater does not suggest that the original evidence *itself* does not support that the object is red. That something seems red makes it plausible that it is red. On the other hand, when I get evidence that an epistemic peer of mine inferred $\neg P$ from the same batch of evidence from which I inferred P , this is evidence that my original FOE supports something other than what I assumed it did.

My account of the distinction is similar to the accounts of other authors. Lasonen-Aarnio (2014: 317) captures the distinction by saying that HOE is evidence that the belief one formed on the basis of one's original FOE was irrational, whereas ordinary undercutting defeaters are not evidence to this effect. This matches my way of putting it since I characterize HOE as evidence that can show that one's belief was not supported by one's original FOE, or, in other words, that one's belief was irrational. Ordinary undercutting defeaters, on the other hand, only have the effect of changing what one's new body of evidence supports, and thus don't imply anything about the reasonableness of forming the relevant belief from one's original FOE. Furthermore, Christensen (2010: 195) suggests that what is characteristic of HOE is that a rational response to it is to "bracket" one's FOE. That is, in order to rationally account for HOE, one needs to ignore one's FOE and revise one's doxastic attitudes on the basis of this. Ordinary undercutting defeat, however, does not demand that one bracket one's original evidence. Rather, if you want to rationally react to it, you need to add it to your original FOE and revise your belief in light of this new, enlarged body of evidence. This is consistent with my account. If, as I suggest, undercutting defeat does not have an effect on what is rational for you to believe by indicating

what your original FOE supports, then it must do so by changing what the body of your FOE supports.³

4.3 Motivating Steadfastness and Conciliationism

In this section, I want to present some of the respective considerations in favour of Conciliationism and Steadfastness. For this purpose, let us look at some examples of misleading HOE. Proponents of Conciliationism often rely on cases like the following to motivate their position.

Logical Reasoning. I've been presented with the following facts: $\forall x(Fx \rightarrow Gx)$ and $\neg(Fa \wedge Ga)$. I infer $\neg Fa$. I then get evidence that I've been slipped a drug which distorts deductive reasoning and makes people in 50% of all cases come to the wrong conclusions.⁴

The relevant bodies of evidence and the target proposition are:

- FOE_{LR} : $\forall x(Fx \rightarrow Gx)$ and $\neg(Fa \wedge Ga)$
- HOE_{LR} : I've been slipped a drug that makes my deductive reasoning 50% reliable
- Target proposition: $\langle \neg Fa \rangle$

Many have the intuition that a rational reaction to receiving HOE_{LR} is to change my belief that $\neg Fa$ to suspension of judgement on whether $\neg Fa$.

There are other kinds of HOE that apparently undermine deductive reasoning. I have already mentioned peer disagreement:

Peer Disagreement. My friend and I have often amused ourselves by solving little math problems in our heads, and comparing our answers. We have strikingly similar track records: we are both very reliable at doing mental maths, and neither is more reliable than the other.

³Feldman (2005: 113) suggests an alternative approach for telling the two apart. I reject this approach for the reasons given by Christensen (2010: 194).

⁴This case is Weatherson's (ms: sec. 1.5) variant of a case described by Christensen (2010: 187). I prefer Weatherson's version since here it's clearer that we have a body of FOE that entails the target proposition, which will be of importance later.

We now engage in this pastime, and I come up with an answer to a problem, 457. I then learn that my friend came up with a different answer, 459.⁵

We're assuming that I am right and that the answer is 457. The relevant bodies of evidence and the target proposition are:

- FOE_{PD} : Details of the mathematical problem
- HOE_{PD} : My peer disagrees with me
- Target proposition: <The answer is 457>, short *457*

Again, it's a widespread intuition that in light of my friend disagreeing with me, I should suspend on *457*. There are also cases in which HOE seems to have an impact on doxastic attitudes that are based on non-deductive reasoning like abductive or inductive reasoning:

Abductive Reasoning. I'm a medical resident. I correctly recognize a patient's symptoms X and come to the conclusion that he has arthritis. Furthermore, that the patient has arthritis is in fact the best explanation for the patient having symptoms X —that is, my diagnosis was not a lucky guess. However, a nurse comes up to tell me that I've been up for 36 hours, and I know that the diagnoses of sleep-deprived doctors can be right just as well as wrong.⁶

The relevant bodies of evidence and the target proposition in this case are:

- FOE_{AR} : Patient has symptoms X
- HOE_{AR} : I have been up for 36 hours and the diagnoses of sleep-deprived doctors can be right just as well as wrong
- Target proposition: <The patient has arthritis>, short *Arthritis*

⁵This is taken verbatim from Lasonen-Aarnio (2014: 315).

⁶This is a variant of a case described by Christensen (2010: 186), too. Horowitz (2014: 719) and Coates (2012: 114) discuss similar cases.

The reasoning process here is inference to the best explanation, i.e., non-deductive. To many it seems that I should give up my belief that *Arthritis*, and suspend judgement on the proposition.

Our intuitions about these cases (or at least the intuitions of many) provide support for Conciliationism. Perhaps the main theoretical consideration Conciliationism draws on is the idea that in order to assess how likely it is that my judgement on the issue in question is correct, e.g., that the answer is 457 or that the patient has *Arthritis*, I have to reason about this *independently* of my judgement on the issue. It would be question-begging to reason that since the FOE in fact supports 457, my judgement is very likely correct (Christensen 2011: 1; Elga 2007: 490). In other words, in order to assess the probability of my judgement being correct, I can only rely on my HOE. In the case of peer disagreement, my HOE is that you and I have come to inconsistent conclusions on the issue. Since, per stipulation, I assume that we're equally reliable, I have to think it is 50% likely that I'm correct. But if I think it's just as likely as not that I'm right, then I better suspend judgement. Similar reasoning applies for other kinds of HOE. If I assign, on the basis of my HOE and independent of my first-order reasoning about whether my evidence in **Logical Reasoning** entails $\neg Fa$, a probability of 50% to my belief that $\neg Fa$ is correct, I should come to suspend judgement on this issue (Christensen 2010: 195-197).

Against this, defenders of steadfast views have presented different arguments. First, Kelly (2005, 2010) accuses conciliatory views of “throwing away evidence”, to put it in Christensen’s (2011: 2) terms. Consider **Peer Disagreement**. I have my FOE, which, as stipulated, supports 457. It supports it in the strongest possible sense: it entails it. I now learn that my peer disagrees with me. Kelly argues that when we want to know how we should react to HOE, we should consider what we should believe in light of the *total* evidence we possess after having gathered the HOE. In Kelly (2005), he claims that besides the HOE that my peer believes that the answer is 459, there is also the HOE that I believe that it is 457. Given that we're equally reliable, I should give equal weight to these bodies of evidence, such that they cancel each other out. If we add to this that my total body of evidence further contains my FOE, which strongly supports 457, it follows that my total body of evidence remains supporting 457. Consequently, I should believe 457.

Weatherson (ms: sec. 2) makes the same argument. He adds that defenders of Conciliationism must therefore reason as follows: if, for example, I judge that 457 in **Peer Disagreement**, this judgement ‘screens off’ my original evidence for 457 and becomes my only evidence in favour of 457. Since my only remaining evidence concerning whether 457 is your judgement that the answer is 459, my evidence overall accordingly supports withholding on whether 457. Weatherson (ms: sec. 4) then goes on to argue against this that the assumption that one’s judgement on a proposition P screens off one’s original evidence on P is false. If it were true, one’s second-order judgement that one has judged P should then also screen off one’s first-order judgement, and one’s third-order judgement that one has judged that one judged P should screen off that second-order judgement, etc. We then end up with an infinite regress to the effect that we don’t have any evidence at all.

Finally, Lasonen-Aarnio (2014) chooses a somewhat different strategy. She claims that in cases where our HOE falsely indicates that we made a mistake in our reasoning, our belief is actually the result of a correct epistemic rule. Take the example of the rule that if one knows P and knows $\langle \text{If } P, \text{ then } Q \rangle$, one should infer Q . Assume that I am in a situation where I know P and $\langle \text{If } P, \text{ then } Q \rangle$, and accordingly come to believe Q . I now gather HOE that modus ponens is incorrect. If HOE stirs us away from what would be the outcome of following a correct epistemic rule, we should rather stick with our belief that Q . One might say that the rule is not correct since precisely because HOE that modus ponens is invalid can undermine an inference from P and $\langle \text{If } P, \text{ then } Q \rangle$ to Q , we shouldn’t always make this inference. Rather our rule should include a *proviso* like ‘...and we don’t have evidence that modus ponens is invalid’. To this, Lasonen-Aarnio (2014: 323) replies that there could also be HOE that undermines beliefs based on such a rule, for example disagreement with a peer over whether there is any evidence against the validity of modus ponens. Of course, one could integrate a further proviso into our rule to deal with such cases, but this would just lead us into an infinite regress.

The literature on HOE, and in particular peer disagreement, is vast, and there are many more arguments brought forward on each side. I won’t discuss them here. The point of this section was just to motivate the two views and give a rough account of what’s respectively appealing about them. They both have

something going for them. So it would be nice if we could reconcile them. In the remainder of this chapter, I show how to do this.

4.4 The Evidence-Ambiguity Theory

4.4.1 The Semantics of the Doxastic ‘Should’

In this section, I give a more precise account of the meaning of ‘should_{FOE}’ and ‘should_{HOE}’ and show how they fit into the simplified account of the semantics of the doxastic ‘should’ that I introduced in chapter 2:

Doxastic Should For any context of utterance c that provides a doxastic ordering source, ‘ S should have doxastic attitude D to P ’ is true at c iff having D to P reflects accurately how much i supports P , where i is the information-state provided by c .

For the following discussion, it is quite relevant what it means for an information-state i to support a target proposition P and how we measure the strength of this relation. As I have explained in chapter 2, I think of this support relation as what we ordinarily understand as evidential support. A strong inferential relation between i and P , be it deductive or abductive, or a high evidential probability of P on i will often instantiate a strong support relation between i and P . I have also pointed out that evidential support is neither reducible to or co-extensive with any one of these inferential or probabilistic relations or their disjunction, as we otherwise end up with the counter-intuitive result that every body of evidence supports every logical truth. Since I cannot solve the question of what the nature of evidential support is, we will have to rely on an intuitive understanding of what the evidential relations are between the information-states and the target propositions that we will consider in the following.

In the remainder of this section, I explore how to fit the distinction between should_{FOE} and should_{HOE} into **Doxastic Should** by discussing two different approaches. I will reject the first and accept the second one.

4.4.2 Approach 1

A first natural thought is that should_{FOE} is relativized to FOE, while should_{HOE} is relativized to both the FOE and the HOE, meaning that, if we model bodies of evidence as sets of possible worlds, it's relativized to $FOE \cap HOE$, the intersection of the two. This fits how Feldman (2005) thinks about how one should react to HOE: it is just more evidence, which is added to the evidence one already possesses (one's FOE). What one should_{HOE} —or for Feldman, simply should —believe is what one should believe in light of this more informed body of evidence.

A problem with this approach, however, is that the support relation between FOE and P will, at least in some cases, also hold between $FOE \cap HOE$ and P . This is problematic where the relevant HOE is meant to undermine FOE. In such scenarios, the same doxastic attitude to P will be licensed (i) by the support relation between FOE and P and (ii) the support relation between $FOE \cap HOE$ and P , where we actually want different doxastic attitudes to be licensed by these relations.

This is most obviously the case where one's FOE entails the target proposition, as in **Logical Reasoning**. $FOE_{LR} \cap HOE_{LR}$ also entails $\neg Fa$. So if it is the entailment relation between FOE_{LR} and $\neg Fa$ that makes it the case that I should_{FOE} believe $\neg Fa$, then it seems that in light of $FOE_{LR} \cap HOE_{LR}$, I should also believe $\neg Fa$.

The problem is not limited to deductive support relations. Christensen (2010: 197) claims that if there exist explanatory connections between FOE and a target proposition P , “[t]hese connections [...] do not depend on any claims about me, and the new information I learn about myself [i.e., the HOE] does not break these connections”. Christensen does not provide any arguments for his claim, so let's consider **Abductive Reasoning** to have a closer look at it.

The relevant bodies of evidence and the target proposition in the case are:

- FOE_{AR} : Patient has symptoms X
- HOE_{AR} : I have been up for 36 hours and the diagnoses of sleep-deprived doctors can be right just as well as wrong
- Target proposition: *Arthritis*

FOE_{AR} supports *Arthritis* in virtue of the fact that *Arthritis* is the best explanation for the patient having symptom *X*. Christensen is certainly right that FOE_{AR} has this support relation to *Arthritis*, irrespective of whether HOE_{AR} is true or not. The fact that I am sleep-deprived seems to be irrelevant to the medical question of whether *Arthritis* is the best explanation for the patient having symptoms *X*.

However, it is not clear that just because HOE_{AR} does not bear on the abductive support relation between FOE_{AR} and *Arthritis*, it is guaranteed that the same support relation holds between FOE_{AR} \cap HOE_{AR} and *Arthritis*. Take a case where you have FOE that all of the 100 swans you have observed so far have been white. This body of FOE gives strong enumerative inductive support to the proposition <All swans are white> (or at least the proposition <most swans are white>). Then you gather more FOE, which consists of an observation of 100 black swans. While the fact that you've observed 100 black swans doesn't make a difference to the support-relation between your previous body of FOE and the proposition that all swans are white, the intersection of this new FOE and your previous FOE does obviously not have the same support relation to the proposition <All swans are white>.

Thus, even if we pool a body of FOE with another body of evidence that does not bear on the evidential relation the original FOE has to the target proposition *P*, their intersection might have a different support relation to *P*. To check which support relation holds between *Arthritis* and FOE_{AR} \cap HOE_{AR}, let's examine whether *Arthritis* is the best explanation for FOE_{AR} \cap HOE_{AR}. Whether it is depends on the alternative possible explanations. For the sake of simplicity, let's assume there is only one other possible disease—lupus—which could have caused symptoms *X*, but that the symptoms are more typical for arthritis than lupus. Therefore, *Arthritis* is a better explanation for symptoms *X* than the proposition that the patient has lupus (*Lupus*). Now, if we add evidence that I was sleep-deprived when concluding *Arthritis* (HOE_{AR}) to my evidence that patient has symptoms *X* (FOE_{AR}), is this joint body of evidence (FOE_{AR} \cap HOE_{AR}) better explained by *Arthritis* or *Lupus*? I suggest we answer this by reasoning in the following piecemeal fashion:

First, neither *Arthritis* nor *Lupus* is a better explanation of HOE_{AR}, i.e., the proposition that I made my diagnosis while I was sleep-deprived and that such

diagnoses can just as well be wrong as right. If anything, neither of them explains this. Second, *Arthritis* explains the FOE_{AR}-part of the joined body of evidence better than *Lupus* does. Combining these two factors, we can conclude that *Arthritis* explains the joint body of evidence better than *Lupus* does. Hence, this joint body of FOE and HOE licenses believing *Arthritis*, rather than suspending judgement on it, just as the original FOE did.

I conclude that in all the three cases described above, the joint body of evidence resulting from adding the HOE to the FOE has the same support relation with the target proposition as the FOE alone has. So, Approach 1 fails.

4.4.3 Approach 2

Approach 2 says that where ‘should’ means should_{FOE}, the information-parameter is filled with the subject’s FOE; and where it means should_{HOE}, the information-parameter is filled with the subject’s HOE. This idea reflects Christensen’s suggestion that in order to respect HOE, we have to bracket our FOE. I think Approach 2 is the way to go for the Evidence-Ambiguity Theory.

Approach 2 gets the cases discussed in section 4.3 right. That I have been slipped a drug that makes my deductive reasoning 50% reliable makes it as likely as not that my belief that $\neg Fa$ is true. Therefore, HOE_{LR} neither supports $\neg Fa$ nor its negation, and I accordingly should suspend judgement on $\neg Fa$ in light of my HOE in **Logical Reasoning**. Similarly, the fact that my friend believes that the answer is 459 and I believe that it is 457 makes it as probable as not that the answer is 457. Thus, in light of my HOE in **Peer Disagreement**, I should suspend judgement on whether the answer is 457 or 459. Finally, in **Abductive Reasoning** I come to know that I made my diagnosis while I was sleep-deprived and that my diagnosis could hence be right just as well as wrong. My HOE in **Abductive Reasoning** accordingly neither favours *Arthritis* nor its negation, and I should_{HOE} suspend judgement on *Arthritis*.

A first objection to Approach 2 might be that it overgeneralizes. Consider **Logical Reasoning**. The relevant bodies of evidence and the target proposition

were:

- FOE_{LR} : $\forall x(Fx \rightarrow Gx)$ and $\neg(Fa \wedge Ga)$
- HOE_{LR} : I've been slipped a drug that makes my deductive reasoning 50% reliable
- Target proposition: $\langle \neg Fa \rangle$

I have claimed that there is no support relation between HOE_{LR} and $\neg Fa$. HOE_{LR} neither supports $\neg Fa$ nor its negation, Fa . It follows from **Doxastic Should** that I should $_{HOE}$ suspend judgement: suspending judgement is the doxastic attitude to $\neg Fa$ that best reflects how much HOE_{LR} supports $\neg Fa$.

Yet something about this explanation seems off. For any information-state i and target proposition P such that i neither supports P nor its negation, it is true that the doxastic attitude that most appropriately reflects this support relation is suspending judgement on P . **Doxastic Should** would therefore predict that for any piece of HOE that has no support relation with a target proposition P , S should (in the sense of should $_{HOE}$) suspend judgement on P . But is this right? Imagine the following case:

Tree. Hannes sees a tree in front of him and comes to believe that there is a tree in front of him. Now he gathers HOE that his mathematical reasoning is distorted. This is the only HOE he possesses.

Hannes' HOE, HOE_{Tree} , has no support relation with the target proposition $\langle \text{There is a tree in front of me} \rangle$, short *Tree*. If we assume that a lack of support is appropriately reflected by suspending judgement, it follows that Hannes should $_{HOE}$ suspend judgement on whether there is a tree in front of him. This is odd. It seems that he should not, in any sense, suspend judgement on whether there is a tree in front of him. He has clear perceptual evidence supporting that there is a tree in front of him and the HOE has no undermining effect on this.

I bite the bullet and maintain that Hannes should $_{HOE}$ suspend judgement on whether there is a tree in front of him. If some evidence supports neither P nor its negation, one should, *in light of this evidence*, support judgement on whether P . I agree, however, that there is something odd about saying that

Hannes should suspend judgement on whether there is a tree in front of him. The reason for this oddness is purely pragmatic: there is no natural context in which one would discuss whether Hannes should believe that there is a tree in front of him and in which this HOE would be the information-state provided by this context. Thus, there is no natural conversation where one is talking about what Hannes should believe in **Tree** and ‘should’ is relativized to the HOE Hannes has in **Tree**. As I will explain in section 4.6, a context’s information-state shifts from a subject’s FOE to their HOE if the HOE undermines or ‘brackets’, to put it as Christensen (2010), their FOE. Since HOE_{Tree} does not undermine Hannes’ FOE supporting *Tree*, this HOE does not become the information-state provided by a normal context in which it is discussed what Hannes should believe. True, I can force ‘should’ to have the relevant sense by saying ‘in light of Hannes’ HOE that his mathematical reasoning is distorted, he should...’, but a context wouldn’t naturally give ‘should’ this sense.

Note that from the truth of ‘Hannes should suspend judgement on *Tree* in light of HOE_{Tree} ’ it does not follow that Hannes has a reason to suspend judgement on *Tree*, which would be absurd. It is the FOE-undermining power of HOE that makes it provide reasons. When I gather HOE_{LR} , this undermines my FOE concerning $\neg Fa$, which gives me reason to suspend judgement on whether $\neg Fa$. HOE_{Tree} , in contrast, does not undermine Hannes’ perceptual evidence concerning *Tree* and hence does not give him reason to suspend judgement on whether *Tree* is true. That HOE_{Tree} can’t undermine Hannes’ perceptual evidence, whereas HOE_{LR} can undermine my FOE concerning $\neg Fa$, is obviously so because HOE_{Tree} is not HOE with respect to *Tree*, but HOE_{LR} is HOE with respect to $\neg Fa$. Whether my deductive reasoning was distorted when I concluded that $\neg Fa$ has something to do with whether $\neg Fa$ is true, whereas whether Hannes’ mathematical reasoning was distorted when he perceived the tree does not seem to have any bearing on whether there is a tree in front of him. For a piece of HOE to be HOE with respect to a target proposition P , it needs to be evidence about FOE with respect to P . Given the account of HOE I presented in section 4.2, this means that a body of evidence is HOE with respect to P iff it is evidence about either (i) what FOE one possesses with respect to P , (ii) what the evidential relation between one’s FOE and P is, or (iii) the reliability of one’s cognitive faculties responsible for processing FOE on P . HOE_{Tree} is accordingly not HOE

with respect to *Tree* since it ticks none of these boxes in relation to *Tree*.

A second objection likely to be made against Approach 2 is that it ignores the norm that one should form one's beliefs based on one's *total* evidence. Both what one should_{FOE} and what one should_{HOE} believe is relativized only to a part of one's total evidence, namely one's FOE and one's HOE. However, it seems to be a platitude that when forming one's beliefs one should take all of one's evidence into account. Let's call this position following Carnap (1950) the *requirement of total evidence*.⁷

For a start, I should point out that a lot of the statements I defend in this thesis are, at first sight, at odds with the requirement of total evidence. For example, in the previous chapter, I have defended the view that one should, in the objective sense of 'should', believe and only believe what is true. This is often not what one's total evidence supports. Furthermore, we sometimes talk about what someone should believe in light of the evidence they had at some point in the past, which is different from their current total evidence.⁸ Finally, in the next chapter I will propose that when you give epistemic advice you talk about what your advisee should believe in light of the collective evidence, i.e., the combined evidence of the two of you, which often goes beyond your advisee's total evidence.

This is not to say that I necessarily reject the requirement of total evidence. Rather, I believe that there are a number of senses of the doxastic 'should' and that the requirement of total evidence does not apply to all of them, as for example the collective or the objective one. Yes, it might be true that in some sense of 'should', one should believe what is supported by one's total evidence, but certainly not in every sense of 'should'. I would argue that should_{HOE} is another sense of the doxastic 'should' that is not tied to the requirement of total evidence. Christensen's (2010) view that HOE brackets one's FOE provides support for this idea. He believes that in response to gathering HOE, one should ignore one's FOE, rather than take it into account.

Kelly (2005, 2010) has criticised Christensen's view precisely because he thinks that it violates the requirement of total evidence. Unlike Christensen, I have the resources to respond to this worry by claiming that I don't disagree with Kelly

⁷The requirement is also endorsed, among others, by Adler (1989), Williamson (2000: 189), and Kelly (2008: 938).

⁸See the case **Nostalgia** in subsection 2.6.2.

that in *some* sense the requirement of total evidence has to be respected. The sense I have in mind is ‘should_{FOE}’. As already mentioned in section 4.3, Kelly (2005) himself suggests that in **Peer Disagreement** the total evidence supports what the FOE supports since the HOE, i.e., that I believe that the answer is 457 and my peer believes that it is 459, outweighs each other.⁹ Furthermore, I have argued in the previous subsection, when discussing Approach 1, that all the bodies of FOE in the cases in section 4.2 do not change the support relations with their target proposition even if the respective bodies of HOE are added to them. In short, if the arguments in the previous subsection are sound, then what one should_{FOE} believe and what one should believe in light of one’s total evidence is co-extensive in the cases discussed.

One might protest that even if following the requirement of total evidence and believing what one should_{FOE} each leads to the same doxastic attitudes, the latter is still a violation of the former since it only requires one to believe what is supported by one’s FOE, not one’s total evidence. The worry seems overly pedantic. If the effects of following two different norms is the same, why would it matter which norm one follows?

In any case, there is a way to fully reconcile Approach 2 with the requirement of total evidence. If one’s FOE and HOE combined have the same support relations with every target proposition as FOE alone has, then we could stipulate that it is not our FOE that the doxastic ‘should’ is sometimes relativised to, but rather our total evidence. That is, I can propose that defenders of Steadfastness are not talking about what one should_{FOE} believe, but about what one should believe in light of one’s total evidence, or short, what one *should*_{TE} believe. On this version of the Evidence-Ambiguity Theory, Steadfastness and Conciliationism are not in conflict because defenders of the former are talking about what one should_{TE} believe, whereas defenders of the latter are talking about what one should_{HOE} believe.

⁹Kelly (2010: 143) weakens this claim. He reasons that since the total body of evidence now contains two pieces of HOE that cancel each other out and thus together neither support that the answer is 457 nor 459, the whole body now supports 457 to a lesser degree than it did before. For this reason, one should lower one’s confidence in 457 somewhat, even though by far not so much that one comes to suspend judgement. Matheson (2009: 274) shows that this argument is flawed. By this reasoning, adding any kind of evidence that doesn’t bear on our target proposition *P* and thus neither supports *P* nor $\neg P$ to our total body of evidence should push us towards suspending on *P*. This is obviously absurd.

One might attack the idea that one’s FOE and one’s total evidence always have the same support relations with any target proposition by pointing out that HOE can also work as FOE. If I get HOE that my FOE supports P isn’t this also evidence for P ? As Worsnip (2015b: 23) puts it, the effects of HOE sometimes “trickle down” to FOE. Thus, if HOE sometimes has FOE-like effects, then, in some cases, one’s total evidence should have different support relations with some target propositions compared to FOE on it’s own. I reply to this that if a body of evidence E , which is HOE with respect to P , is also FOE with respect to P , then E is not only part of our HOE, but also of our FOE. Thus, combining FOE and HOE and thus ‘adding’ E to FOE does not make a difference in such cases since E would be part of one’s FOE already anyway.

To conclude, the requirement of total evidence is already respected by the Evidence-Ambiguity Theory in that what one should_{FOE} believe is co-extensive with what one should_{TE} believe in the cases discussed in section 4.3. If this does not suffice for a proponent of the requirement of total evidence, I could adjust the Evidence-Ambiguity Theory such that defenders of Steadfastness are talking about what one should_{TE} believe, not what one should_{FOE} believe. In the following, however, I will put this issue aside and return to speaking of what one should_{FOE} believe.

4.5 So What Should I Believe?

I have shown that the distinction between should_{FOE} and should_{HOE} can be integrated into my information-sensitive contextualist semantics for the doxastic ‘should’. Now, a main concern with contextualist theories about ‘should’ is that they seem to dodge the relevant question. After all, we got into the debate about peer disagreement, or HOE in general, hoping to find an answer to the question of what one should believe in such scenarios. The Evidence-Ambiguity Theory doesn’t answer this question, but rejects it as ambiguous: is the question what one should believe in light of one’s FOE or is it what one should believe in light of one’s HOE? To this it can be objected that there *seems* to be a clear, unambiguous question. In this and the next section, I try to resolve this apparent problem for the Evidence-Ambiguity Theory.

First of all, that modals like ‘should’ are context-sensitive *in some way* is

widely accepted among semanticists, and it is thus pretty much uncontroversial that the question is ambiguous *to some degree* if taken out of context. Since ‘should’ can be epistemic or deontic, ‘What should I believe?’ can either mean ‘Which belief do I probably have?’ or ‘Which belief is required by the norms?’ Furthermore, as ‘should’ can express different norms, the latter question can, for example, be specified as either ‘Which belief is required by the practical norms?’ or ‘Which belief is required by the epistemic norms?’¹⁰

Of course, the opponent of the Evidence-Ambiguity Theory might say that this is all good and well, but that the question we have in mind with respect to HOE is ‘What should I believe *in the doxastic sense?*’ and that this question is unambiguous. Against this, I would point to the distinction between the subjective and the objective doxastic ‘should’ that I defended in chapter 3. If we accept that ‘What should I believe?’ can be relativized to different information-states, i.e., sometimes to the facts, sometimes to the subject’s evidence, it is not that much of a step further to posit that it can be relativized to more fine-grained information-states, such as the subject’s FOE and HOE. Certainly, it does not follow from the subjective/objective distinction that there must be a distinction between what one should_{FOE} and should_{HOE} believe, but it puts pressure on the idea that we can, *prima facie*, expect the doxastic question to be unambiguous.

Gibbons (2013) rejects ambiguity theories in normative epistemology as missing the point. As I have mentioned in chapter 3, he thinks that the respective plausibility of the justification norm and the truth norm create a deep philosophical puzzle. Because there are unjustified false beliefs, we can sometimes satisfy one norm, but not the other. Now, what should we do in such a scenario? Are we required to change our false belief or are we rather required to stick with it? While he acknowledges that ‘should’ in English might be ambiguous between a subjective and an objective sense, this does not really resolve the puzzle in his eyes (Gibbons 2013: 21). The central question is whether the subjective or the objective doxastic ‘should’ figures in the question ‘What should I believe?’ that leads us in theoretical reasoning.¹¹

A potential complaint with the Evidence-Ambiguity Theory could accordingly be that it does not tell us whether the deliberating agent is concerned with what

¹⁰For more details on the general context-sensitivity of modals, see chapter 2.

¹¹For more details on this, see section 3.5.

she should_{FOE} or with what she should_{HOE} believe. I reply to this worry in the next section by giving a contextualist model for theoretical deliberation.

4.6 A Contextualist Model for Theoretical Deliberation

In general, the contextualist's position on what the meaning of 'What should I believe?' is, is obvious: it is whatever meaning this question has in the agent's context. A comparison: the meaning of the context-sensitive question 'Who am I?' that a soul-searching agent is asking herself is determined by what 'I' refers to in the context of the agent: herself. Similarly, the meaning of the agent's question 'What should I believe?' is determined by the information-state salient in the agent's context of deliberation.

When an agent starts out to privately deliberate about whether P , they focus on their FOE concerning P . Private deliberation means that they are deliberating on their own, and not as a member of a group that is inquiring together.¹² The question they are asking themselves is 'What should I believe in light of my FOE?' and the information-state provided by this context is accordingly their FOE. In some cases where an agent is in the process of answering this *FOE-question* by weighing their FOE with respect to P , the agent gathers HOE that supports another doxastic attitude to P than their FOE does. In this case, the agent is not anymore in the position of answering the FOE-question. Therefore, the agent turns to a question, getting the right answer to which is second-best to getting the right answer to the FOE-question: the question of what they should believe in light of their HOE (the *HOE-question*).¹³ Let me explain this in more detail.

Christensen's (2010) talk of a bracketing of the FOE by HOE is a way to explain why HOE sometimes puts agents in positions in which they can't answer the FOE-question anymore. If a deliberating agent becomes aware of HOE undermining their FOE, the FOE gets 'disqualified' from being considered in their deliberation. When I gather HOE in **Logical Reasoning** that I have been slipped a drug that

¹²I will cover the practice of collective deliberation in chapters 5 and 6.

¹³Sometimes, an agent might gather HOE with respect to P , *before* or *at the same time*, rather than after, they gather the FOE with respect to P undermined by this HOE. In such cases, they do not turn from the FOE-question to the HOE-question, but go straight for the HOE-question.

distorts deductive reasoning, I am not competent anymore to answer, based on my FOE, whether I should believe $\neg Fa$ or not. That I am not competent to answer this question might be due to the fact that I don't know what my FOE supports after I gathered the relevant HOE. Once I am in this position, the information-state of the context shifts to the HOE. To some degree, this also makes sense if we consider which metasemantic principle might be at work here. As explained in chapter 2, speaker intentions are a very plausible candidate for determining the contextual parameters of modals, such as the information-state. When I can't answer the FOE-question anymore, I intend to reason about what I should believe about this in light of my newly gathered HOE instead, which is second-best to believing what I should_{FOE} believe. While context-sensitive semantics have been developed mostly with the intention to model context-sensitive expressions in language, their basic principles arguably also apply to context-sensitive *thoughts*. When I think 'I am hungry' and you think 'I am hungry', too, we are thinking two different propositions since 'I' in my thought refers to me and 'I' in your thought refers to you. Similarly, it seems plausible that my deliberating question changes from 'What should_{FOE} I believe?' to 'What should_{HOE} I believe?' when I realize that I can't answer the FOE-question anymore and therefore intend to deliberate about what I should believe in light of my HOE.

The other two candidates for determining a contextually provided information-state that I mentioned are the pragmatic purpose of an utterance and conversational salience. As just hinted at, private deliberation is done in thought and not a communicative act. Therefore, we can presumably ignore these factors here. However, if 'thinker' intentions are the only factor that could determine the sense of 'should' in our thoughts, the following interesting question arises: what happens when someone gathers HOE that undermines their FOE, but they nonetheless intend to think about what they should_{FOE} believe, or what happens when they don't gather such HOE, but intend to think about what they should_{HOE} believe? Contrary to what I have suggested in the previous paragraph, it seems at least not implausible that in such scenarios the information-state does not shift in accordance with whether the subject's FOE is undermined or not, but with the subject's intentions. I can't go into too much detail here and settle which of the two options is correct. All that matters in order to avoid the worry that the deliberating question can't be guiding on a contextualist semantics is that there

will be *some* kind of mechanism determining which sense the deliberating question carries.

Why is correctly answering the HOE-question second-best to correctly answering the FOE-question? Motivation for this claim can be found in Lasonen-Aarnio (ms, 2014) and Schoenfield (2012, ms). Even though they both think that the *rational* thing to do in cases where one gathers misleading HOE is to be steadfast, they claim that it is *reasonable* to revise (Lasonen-Aarnio 2014: 343; Schoenfield 2012: 209). They think it is reasonable—even though not epistemically rational—since it is what someone should do who has the aim of conforming to the norm of believing what one’s FOE supports (Lasonen-Aarnio ms). In other words, they say that following one’s HOE is a *means* to conform with what one should_{FOE} believe.

While I think they are on the right track, this can’t be quite right. If on the basis of my HOE, I know that my FOE either supports P or $\neg P$, but I have no idea which, and therefore suspend judgement on P , this is neither a means to conform with what I should_{FOE} believe nor the result of planning to do so. After all, I *know* that suspending judgement on P is *not* what is supported by my FOE. If I perform action A and I know that doing A can’t make me achieve my goal z , doing A is neither a means to nor the result of planning to achieve z .

I want to put the idea that believing what you should_{HOE} believe is second-best to believing what you should_{FOE} believe in a somewhat different manner. It is a common idea in epistemology that the truth norm for belief is fundamental to other norms like the justification or rationality norm in that they are derived from the truth norm.¹⁴ Accuracy-first approaches to epistemology¹⁵ give a good account of this derivative relation. The truth norm implies a measurement of credences in terms of their accuracy, that is, their closeness to truth. The higher a credence in a true proposition, the more accurate it is. Now, answering the FOE-question tells us which credence to P is the one with the highest *expected* accuracy value. Very roughly, the higher the probability of P on my FOE, the higher will be the credence to P with the maximal expected accuracy (Greaves and Wallace 2006).

In an analogous fashion, we can construe believing what one should_{HOE} believe

¹⁴See Joyce (1998) and Wedgwood (2002).

¹⁵I take this label from Schoenfield (ms).

as second-best to believing what one should_{FOE} believe because it maximizes the expected value of what the norm articulating what one should_{FOE} believe promotes: expected accuracy. Let's call this maximized value your *expected expected accuracy*.

It's beyond the scope of this thesis to go far into the details, but I hope the following example can illustrate what I have in mind. Imagine that on my HOE, there is a 50% chance that my FOE fully supports P and a 50% chance that it fully supports $\neg P$. Intuitively, I should suspend judgement on P (in some sense). In fact, my FOE fully supports P and the expected accuracy on my FOE of believing P is 1, and the expected accuracy of believing $\neg P$ is 0. If my evidence fully supported $\neg P$, the expected accuracy on it of believing P would be 0, and the expected accuracy of believing $\neg P$ would be 1. The expected expected accuracy of a doxastic attitude D relative to a body of FOE F and a body of HOE H is calculated as follows:

Expected Expected Accuracy $EEA_{F,H}(D) = \sum_k (Pr_H(k) \times EA_k(D))$, where $k \in K$ and K is a partition of ways F could be like, $Pr_H(k)$ is the probability of k obtaining on H , and $EA_k(D)$ the expected accuracy of D on k .

In the given example, there are two members of k :

- k_1 : Your FOE fully supports P
- k_2 : Your FOE fully supports $\neg P$

The expected expected accuracy of believing P is the average of the expected accuracy (EA) of believing P given the different ways my FOE could be, weighted by the probability on my HOE of my FOE being the respective ways. Given that there is a 50% probability for k_1 and k_2 respectively, $EEA(\text{Believing } P) = 0.5$ and $EEA(\text{Believing } \neg P) = 0.5$. Now, where believing P and believing $\neg P$ are 'equally good' from an epistemic perspective—in our case in terms of expected expected value—suspending judgement is the best option. Thus, suspension of judgement will be the (coarse-grained) doxastic attitude maximizing the expected expected value.

It is not important that this particular way of cashing out the idea that believing what one should_{HOE} is second-best to believing what one should_{FOE}

works. Maybe Lasonen-Aarnio's (2014), Schoenfield's (2012), or another account that I have not considered is better for this. All that matters is that there is *some* account of the idea that works. Given that the idea is quite intuitive, I am hopeful that there is some such account.

In the remainder of this section, I will address a number of objections to my contextualist model of doxastic deliberation.

A possible objection to my story about the context-shift from the FOE-question to the HOE-question is that it is controversial to assume that epistemic norms are in a hierarchy relative to the epistemic goal of having true beliefs. For example, when we discussed blindspot propositions in chapter 3.5, we saw that it is at least tricky to formulate the truth norm in a general enough manner so that other norms can be derived from it.¹⁶ Furthermore, one might reject the idea that truth is the epistemic *summum bonum*, and claim that rather knowledge¹⁷ or understanding¹⁸ is. To both these points I could reply that for my purposes it is sufficient if maximizing expected accuracy is second-best to maximizing truth. This can be the case irrespective of whether the importance of maximizing expected accuracy is again derived from the truth norm, i.e., the value of accuracy. Third, however, one might dislike the idea that norms stand in derivative relations to each other altogether, maybe because the proximity of this idea to epistemic instrumentalism. It is important to notice therefore that the claim that we move from the FOE-question to the HOE-question when facing undermining HOE does not rest on the assumption that there is a truth-orientated hierarchy among epistemic norms. This is only the most plausible explanation I could think of for this shift. All this claim requires is that there is *some* reason why privately deliberating agents turn to their HOE if their FOE has become 'bracketed' or 'undermined'. Another explanation could be that such agents simply need to orientate themselves by *something* in their belief-forming business and that their HOE is just the only plausible candidate left after their FOE has been undermined. This story would need to work out what makes HOE the only plausible option. In any case, the general point is that if the truth-hierarchical story won't work, I am confident there will be another one that can convincingly

¹⁶Wedgwood (2013) proposes a solution to this problem.

¹⁷See, for example, Bird (2007), Sutton (2007) and Williamson (forthcoming).

¹⁸See, for example, Kvanvig (2003).

explain why when their FOE gets discredited, privately deliberating agents turn to their HOE.

This last point shows that a lot in my story hinges on Christensen’s (2010) idea that misleading HOE brackets FOE. As already mentioned in section 4.3 above, Weatherson (ms: sec. 1.5) claims that this idea relies on the mistaken assumption that judgements screen evidence. According to Weatherson, we cannot say that misleading HOE undermines or brackets an agent’s original *evidence* concerning a target proposition P , i.e., their FOE concerning P . At best, we can say that it brackets the *judgement* the relevant agent formed on the basis of this FOE. Christensen therefore has to assume that the evidence that gets bracketed by HOE in **Logical Reasoning** is my belief that $\neg Fa$, and not FOE_{LR} , or so Weatherson argues. He thinks Christensen must make this assumption since the misleading HOE, HOE_{LR} , indicates that my belief that $\neg Fa$ is unreliable, and not that the evidence FOE_{LR} from which I’ve formed it is unreliable. But if one’s evidence in such cases is one’s first-order judgement, and not the evidence it was based on, we get into the infinite regress described in section 4.3.

I maintain that the HOE in the cases we have discussed brackets the respective bodies of FOE themselves, and not the judgements formed from them. On my interpretation, that one’s HOE brackets one’s FOE does not mean that one’s HOE shows that one’s FOE is unreliable. Rather, when a deliberator’s HOE brackets their FOE, this means that, from the perspective of this deliberator, their HOE has made their FOE ‘unusable’. A deliberator who has gathered misleading HOE with respect to a target proposition P cannot reasonably—to use Lasonen-Aarnio’s (ms, 2014) and Schoenfield’s (2012, ms) terminology—form judgements about P on the basis of their FOE pertaining to P . It might not be reasonable because ignoring HOE makes one, on average, worse-off truth-wise. For readers who don’t like such a truth-orientated picture of reasonableness, I have to posit that it is a brute fact that ignoring HOE is unreasonable; a fact that we pick up on in our intuitive reactions to cases like **Logical Reasoning**. One might think that this amounts to begging the question against defenders of Steadfastness like Weatherson (ms). However, since I do not deny Steadfastness, but agree that it accurately captures how one should_{FOE} react to misleading HOE, I can’t be guilty of begging the question against Steadfastness.

This leads me to a final potential worry. One might wonder whether my

account still leaves room for the idea that Steadfastness is right in some way, given that I've argued that any agent who is presented with misleading HOE should turn away from the question of what they should_{FOE} believe. However, this doesn't imply in any way that it's incorrect for a third person to state that such an agent should remain steadfast. What brings it about that deliberating agents should turn away from this question is that they have misleading HOE. A third person judging this agent, however, can ignore this HOE if they know it to be misleading and talk about what this agent should_{FOE} believe. Schoenfield (2012: 209) suggests that in ordinary discourse discussing whether another agent is believing what their FOE supports can be important for us in order to evaluate whether this agent is someone we can defer to. The reason is that we want to defer to agents who believe what they should_{FOE} rather than to those who believe what they should_{HOE} since the former agents' ratio of true beliefs will, *on average*, be higher (Schoenfield 2012: 215). Pointing out that someone is not believing what they should_{FOE} could hence be used to mark individuals whom one should not defer to.

It is this external perspective we take when we state, in line with Steadfastness, that agents in cases like **Deductive Reasoning**, **Peer Disagreement**, or **Abductive Reasoning** should maintain their beliefs. And it is the internal perspective of these agents we take when we state, in line with Conciliationism, that they should suspend judgement.

4.7 Dilemma or Ambiguity?

I have argued that the distinction between should_{FOE} and should_{HOE} can be accommodated by Doxastic Contextualism. The upshot was that we don't have to choose between Conciliationism and Steadfastness. Since both views have something going for them, my 'third way' seems therefore preferable to choosing between them. Of course, there might be a significant disadvantage to the Evidence-Ambiguity Theory that outweighs the advantage of having one's cake and eating it, too. I have tried to diffuse worries about one potential disadvantage in sections 4.5 and 4.6.

There is a 'fourth way' that I have not mentioned yet. Instead of claiming that Steadfastness and Conciliationism are both correct and don't conflict, one can

also suggest the view that they are both correct, but actually *do* conflict. Worsnip (2015b: 33) thinks that Christensen (2007, 2010, 2013) should be interpreted in this way. Exegesis of Christensen is not too important here, as the view itself certainly merits to be taken serious.

Worsnip (2015b) uses the example of moral and prudential reasons to illustrate the view. What we morally and what we prudentially should do can come apart. One can interpret this in two ways. On one view, the two kinds of reasons are commensurable and can be weighed against each other such that there is something one should do all-things-considered. For example, one could think that the moral always trumps the prudential. Worsnip (2015b) likens this to (putatively) Christensen's view that there is both a demand on us to follow our FOE and to follow our HOE. On this picture, our FOE and our HOE pull us in opposite directions. There is something we should believe all-things-considered, but it is unfortunate that we can't accommodate the demands put on us both by our FOE and our HOE, just as it is unfortunate when we cannot be both prudent and moral. In other words, misleading HOE places us in a dilemma. Let's call this the *commensurability position*.

The Evidence-Ambiguity Theory is rather like the position that moral and prudential reasons are incommensurable and from completely different normative domains, which do not compete with each other.¹⁹ What we should_{FOE} believe and what we should_{HOE} believe does not compete with each other; they just constitute two different perspectives. Accordingly, there is no dilemma when one cannot both believe what one should_{FOE} and what one should_{HOE}. Let's call this the *incommensurability position*.

Worsnip (2015b) himself defends a version of the incommensurability position. He rightfully points out that its advantage is that it avoids dilemmas (ibid.: 35), which leaves us with the uncomfortable situation that whatever we do, we're violating one of the demands placed on us. Furthermore, it is conservative with respect to the axioms of standard deontic logic. According to the latter, if we should do A and we should do $\neg A$, we should do A and $\neg A$, which is impossible. The incommensurability position resolves this by making clear that we should do A in some sense of 'should' and should do $\neg A$ in some other sense, which does not entail that we should (in one sense) A and $\neg A$ (ibid.: 36).

¹⁹This position is often attributed to Sidgwick, Henry (1874).

I agree with Worsnip (2015b) on this and gladly take his points on for the purpose of defending the Evidence-Ambiguity Theory against the commensurability position. However, we disagree on what the different kinds of normative domains are that are at work in cases where there is misleading HOE. First of all, he is only concerned with cases where one's FOE supports having a doxastic attitude D to a proposition P and one's HOE supports believing that one's FOE does *not* support having D to P (ibid.: 5). He argues that in such scenarios, we have epistemic reasons to have D to P and at the same time epistemic reasons to believe that our evidence does not support having D to P , but that rationality requires us not to have this combination of attitudes as they are incoherent. This conflict between what our epistemic reasons demand and what rationality requires is, according to him, only apparent since reasons and rationality are different kinds of normative things that do not actually compete with each other (ibid.: 34).

On the Evidence-Ambiguity Theory, on the other hand, the two domains are (i) what one should believe in light of one's FOE and (ii) what one should believe in light of one's HOE. One might think that what one should_{FOE} believe is what one has epistemic reasons to believe and what one should_{HOE} believe is what rationality requires, and that Worsnip's and my theory are therefore actually identical. However, they differ in two important aspects. First, as mentioned, Worsnip's theory is only limited to cases where one's FOE supports having a doxastic attitude D to a proposition P and one's HOE supports believing that one's FOE actually does not support this attitude. This already excludes **Logical Reasoning** from the list of cases his theory can be applied to. Remember that in **Logical Reasoning** my FOE supports believing $\neg Fa$. However, my HOE does not support believing that my FOE supports not believing $\neg Fa$, as there is a 50% chance on my HOE that my belief that $\neg Fa$ is entailed by my FOE and a 50% that it is not. My HOE in **Logical Reasoning** does therefore support suspending judgement on what my FOE supports. It follows that **Logical Reasoning** is not the kind of case Worsnip's theory covers, which makes the Evidence-Ambiguity Theory more widely applicable. Second, Worsnip concurs with Broome (1999) that rationality requirements are wide scope such that they only demand of us that we don't have certain combinations of attitudes, but not how to resolve them. The Evidence-Ambiguity Theory, on the contrary, allows that HOE requires

specific doxastic attitudes, namely what we should_{HOE} believe, rather than just forbids combinations of attitudes. This seems to be exactly what is needed given the intuitions we have about the three cases I mentioned in section 4.3, in each of which HOE appeared to demand suspending judgement.

4.8 Conclusion

We can have our cake and eat it, too, in the higher-order debate: we don't have to choose between Conciliationism and Steadfastness. Both views are right in some way. Conciliationism is right about what we should believe in light of our HOE, while Steadfastness is right about what we should believe in light of our FOE. I have defended this Evidence-Ambiguity Theory against the worry that it cannot provide us with an answer to the question of agents deliberating about what to believe. Depending on the respective context of the deliberating agent, a specific information-state is provided by the context, determining which particular meaning the agent's question 'What should I believe?' carries. Finally, I have argued with the help of Worsnip (2015b) that my Evidence-Ambiguity Theory is to be preferred to views according to which we are in a rational dilemma when FOE and HOE conflict. In conclusion, I hope to have presented a further case where highlighting the information-sensitivity of the doxastic 'should' can help us resolve a debate in epistemology.

Chapter 5

The Collective ‘Should’ and Epistemic Advice

5.1 Introduction

In chapters 3 and 4, I have provided evidence for Doxastic Contextualism. I have argued that the doxastic ‘should’ has a objective and a subjective sense, and that it can be further relativised to the subject’s first-order evidence and the subject’s higher-order evidence. In this chapter, I will propose that the doxastic ‘should’ can be relativized to another information-state: the collective evidence, that is, the joint evidence of the speaker and their addressee.

My thesis is that the doxastic ‘should’ carries this *collective* sense in sentences that are uttered for the purpose of giving epistemic advice (epistemic-advice sentences). While there is no account of epistemic advice in the literature, practical advice has received quite a bit of attention, for example from Hobbes (1651/2009) and Searle (1969). Based on their and others’ work, I provide a characterisation of the practice of giving epistemic advice. I suggest that giving epistemic advice is the attempt to help someone who is deliberating about whether a proposition P holds with their epistemic goal of adopting the best-informed doxastic attitude to P possible.

Not only have philosophers written about the nature of practical advice, there is also some recent research on the truth-conditions of sentences uttered for the purpose of giving practical advice (Björnsson and Finlay 2010; Kolodny and

MacFarlane 2010; Dowell 2013). Björnsson and Finlay (2010), at least, think that their account also applies to epistemic-advice sentences.¹ However, neither they nor anyone else in the literature has discussed the specific issues relevant for epistemic advice. The present chapter therefore fills a gap in the literature.

After clarifying what I take epistemic advice to be in section 5.2, I go through several candidates for the truth-conditions of epistemic-advice sentences, that is, through different proposals concerning what the sense of the doxastic ‘should’ in epistemic-advice sentences is. In chapter 3, we have looked at the objective and subjective doxastic ‘should’. I argue that epistemic-advice sentences employ neither of these two (sections 5.3 and 5.4). However, there are further information-states, other than the facts and the subject’s evidence, that theorists could propose as the information-state that the doxastic ‘should’ is relativized to when occurring in epistemic-advice sentences. Other possible candidates are the subject’s—i.e., the advisee’s—available evidence and the speaker’s—i.e. the advisor’s—evidence. I will reject these proposals, too (sections 5.5 and 5.6). Instead, I put forward **Collectivism**. According to this view, epistemic-advice sentences like ‘You should have D to P ’ are true iff D to P is supported by the collective evidence, the combined evidence of the advisor and the advisee (section 5.7).

Figuring out under which conditions epistemic advice is correct helps us to better understand the practice of epistemic advice. Besides this more general theoretical benefit of giving an account of the truth-conditions of epistemic-advice sentences, **Collectivism** also has a more specific advantage: it solves a problem in normative epistemology. Turri (2012) presents a puzzle about withholding. He considers various solutions to it, but in the end rejects all of them and concludes that the puzzle seems to be unsolvable. In section 5.8, I show how **Collectivism** can solve it.

5.2 What Is Epistemic Advice?

To state the obvious first: epistemic advice is a kind of advice. Advice can come in different linguistic forms. One can explicitly say ‘*I advise you* to exercise more’ (Hinchman 2005: 359). Furthermore, one can use the imperatival form, as in ‘Exercise more!’ (Hobbes 1651/2009: Ch. xxv, 3); the subjunctive form, as in ‘If

¹Personal conversation with Steve Finlay.

I were you, I'd exercise more'; or weak deontic necessity modals like 'should' or 'ought', as in 'You should exercise more'.²

My focus here is on this last form. I am looking for the truth-conditions of sentences of the form 'You should have doxastic attitude D to proposition P ' that are used for the purpose of epistemic advice, i.e., of epistemic-advice sentences. However, these truth-conditions will arguably give the conditions for when it is correct to give the same advice in one of the other three forms I've described. That is, I suggest that it is correct for me to utter 'Believe P !' as a piece of advice (instead of, say, a command) if and only if it is true for me to utter 'You should believe P ' as a piece of advice. This does not mean that the truth-conditions of 'You should believe P ' are the truth-conditions of the other forms of giving the same advice. For example, imperatives like 'Believe P !' are usually taken to have no truth-conditions at all. Furthermore, since I'm fallible with respect to whether I'm giving advice, the self-report 'I advise you to believe P ' has different truth-conditions than 'You should believe P '. In conclusion, detecting the truth-conditions of epistemic advice of the form 'You should believe P ' is interesting beyond merely finding out these truth-conditions, since it also provides us with the conditions for correctness (but not truth) for the same advice in a different form.

Epistemic advice needs to be distinguished from practical advice. As a first approximation, practical advice is advice on what one should do, whereas epistemic advice concerns which doxastic attitude one should adopt. Many acknowledge, however, that there are practical reasons for belief.³ The distinction between epistemic and practical advice can therefore not be drawn simply in virtue of the kind of thing—an action or a doxastic attitude—that the advisor recommends to do or to adopt.

What clearly separates the two kinds of advice is that they concern different sorts of normativity. Practical advice is advice on what one should do (or believe) according to practical norms. Epistemic advice, on the other hand, is about what one should believe (or do) according to epistemic norms. According to my characterization of epistemic norms in chapter 1, they concern which doxastic

²This list is not intended to be exhaustive; there might be other linguistic forms. There certainly are furthermore non-linguistic ways of giving advice, as when I am pointing in a certain direction after you've asked me where to go.

³See, for example, Foley (1987: 210) and Papineau (2013: 68).

attitude one should adopt in light of considerations such as truth, evidence, knowledge, or cognitive reliability. Epistemic advice is formulated with the help of the doxastic ‘should’, which expresses epistemic norms, whereas practical advice is formulated with the help of the practical ‘should’, which expresses practical norms.

There are a number of philosophical accounts of advice. They all focus on practical advice, but I think they point out features that epistemic advice shares, too. The *locus classicus* on advice is Hobbes (1651/2009: ch. xxv). He singles out a property of advice that many have agreed on: advising is done with the intention of *benefiting* the advisee. Similarly, Searle (1969: 67) writes that an essential feature of advice is that it is “an undertaking to the effect that” what the advisee is advised to do is in the advisee’s “best interest”. Furthermore, Hinchman (2005: 359) describes prudential advice as being given from the perspective of the advisee’s self-interest.

Wiland (2000: 5) expresses doubts about this, arguing that advice is not only given in order to promote the advisee’s self-interest, but also in order to promote other ends the advisee has. Now, we can incorporate this into Hobbes’ view that advice aims to benefit the advisee by using a wider reading of ‘benefit’, according to which anything that promotes one of my aims, be that a self-interested aim (in the narrow sense of ‘self-interested’) or an altruistic aim, counts as benefiting me. Maybe this fix can’t embrace all kinds of advice. Intuitively, there is such a thing as moral advice, but plausibly what one morally should do sometimes does not promote any of one’s goals. This is most clearly the case when we’re giving moral advice to a sociopath. On the other hand, precisely because a sociopath doesn’t share the goal of doing what’s moral, we might hesitate to call it ‘moral advice’ when we tell a sociopath that he shouldn’t murder innocent people. In any case, I think we can at least say that advice aims at promoting some kind of goal, whether or not that goal is one the advisee actually has or not.

A further important aspect of advice is that it is second-personal. That is, it is about a person or group other than the speaker. I can’t advise myself (at least not the same time-slice of myself). The standard way to give epistemic advice is thus by uttering sentences of the form ‘You should adopt doxastic attitude *D* to *P*’. But I don’t need to use the second-person pronoun. If a philosophy major asks me for advice whether she should take the intro to logic course or the aesthetics

course, and I say ‘Every philosophy major should take at least one logic course’, I have advised her without using the term ‘you’. Furthermore, advice-sentences cannot just be about any other person than the speaker, they must be about the person (or persons) who is the addressee of the utterance. In the literature on deontic modals, eavesdropping cases are discussed, that is, cases where a subject can’t hear the speaker and the speaker talks about what the subject should do in light of her, the speaker’s evidence.⁴ Björnsson and Finlay (2010: 22) describe such utterances as “simulated advice”. This aptly implies that they are not actual advice, precisely because the speaker isn’t addressing the subject. I won’t be concerned with such utterances here.

Let’s take stock. Epistemic advice is given in terms of sentences that express epistemic normativity, it is about the person(s) who is (are) the utterance’s addressee(s), and has the purpose of helping the advisee(s) with respect to an epistemic goal. The satisfaction of all these conditions is still not sufficient for epistemic advice. Remember the following case from section 2.6.2:

Nostalgia. Renaud looks back at a situation at time t , when his evidence clearly supported disbelieving Q . We both know now that this evidence was misleading, and that Q is true. I say to Renaud:

- (1) Don’t beat yourself up. Disbelieving Q is what you should have done at the time.

(1) expresses epistemic normativity and is about Renaud, the addressee of the utterance. One could also say that it’s done for the purpose of helping Renaud with respect to an epistemic goal. As I suggested in chapter 4, we sometimes judge whether someone believes what they should in light of their evidence to evaluate their epistemic reasonableness. Telling you that you believed what you should have believed in light of your evidence at t can be seen as an effort to stabilize a disposition in you to believe in accordance with your evidence, that is, a disposition to be epistemically reasonable.⁵ In other words, uttering (1) is a tiny contribution towards you achieving an epistemic goal, that is, being epistemically reasonable, in the future.

⁴See Björnsson and Finlay (2010: 18), Dowell (2013: 150).

⁵For a paper-length defence of the idea that epistemic-normative talk is used to promote certain belief-forming processes in the members of our community, see Dogramaci (2012).

Admittedly, someone who is giving epistemic advice seems to try to help their addressee to achieve an epistemic goal in a somewhat more straightforward way than I do when uttering (1) in your presence. Being at a loss to single out in which way it is more direct, I'd like to suggest a different way to distinguish between evaluating someone's epistemic reasonableness and giving epistemic advice.

Björnsson and Finlay (2010) and Dowell (2013) put forward the idea that practical advice aims to help practically deliberating agents. Analogously, it seems right to say that epistemic advice tries to help agents who are engaged in doxastic deliberation about whether *P*. Now, in chapters 3 and 4, I have claimed that doxastic deliberators are contemplating about what they should believe in light of their evidence, i.e., in the subjective sense of 'should', sometimes contemplating what they should believe in light of their first-order evidence, sometimes what they should believe in light of their higher-order evidence. This would suggest that the epistemic goal that advisors are trying to assist with is the advisees' epistemic goal to believe about *P* what's supported by their—the advisees'—evidence.

However, on reflection this turns out to be wrong. If I am advising you on what you should believe about whether *P* and I have better evidence concerning whether *P* than you have, then I wouldn't help you if I told you what you should believe in light of *your* evidence. Rather, you would be interested in what my evidence says. Thus, the epistemic goal that epistemic advisors help you with cannot be to believe what's supported by your evidence, otherwise a typical advisor would ignore their own evidence.

There are two ways to interpret this. First, one could conclude that I was wrong in chapters 3 and 4 that doxastic deliberators are wondering about what's supported by their evidence. There is some support for this if we think of deliberation not just as interpreting one's evidence, but also as gathering new evidence. When you're inquiring about when the train to Stirling leaves, you'll not only rely on what your evidence, e.g., your memory, tells you about this, but also look for new evidence, e.g., the train schedule. But we need to distinguish between two different goals of deliberating agents. The goal of inquiry is arguably something like believing what's supported by the best available evidence. We need to distinguish this from the aim of believing what's supported by my current evidence. Feldman (2000) discusses the case of someone who has modest evidence that ginkgo supplements improve one's memory and then sees the cover of a

scientific magazine with the headline ‘The health effects of ginkgo’. She doesn’t know that the article inside says that ginkgo doesn’t have a positive effect on one’s memory. The aim of her inquiry into the effectiveness of ginkgo is to believe what the best evidence on ginkgo supports. Does this mean that right now, before having read the article, she should believe that ginkgo doesn’t boost one’s memory? Feldman (2000: 688) writes: “Until [she’s] read the article, it would be bizarre for [her] to stop believing what’s supported by [her] old evidence”. He seems right. In at least some sense of ‘should’, she shouldn’t stop believing that ginkgo helps one’s memory before she’s read the article. Let’s call this process of forming one’s beliefs in response to the evidence you’ve got *private doxastic deliberation*.

Thus, the view that private doxastic deliberators are guided by the question of what they subjectively should believe is not mistaken. This leaves us with the second option, which is to assume that advisors assist their advisees with another epistemic goal than the one they have as private doxastic deliberators. As the consideration that advisors go beyond their advisees’ evidence suggests, receiving advice is about tapping into the knowledge and experience of others. Corresponding to this, *giving* advice is ‘presumptuous’ (Hinchman 2005: 358). If I tell you what you should do, I seem to think that I know better or at least as well as you what needs to be done. For this reason, the epistemic goal of advisees is not to have the doxastic attitude licensed by an information-state that is limited to their own evidence. They want to profit from their advisor’s evidence. An intuitive way to put this is that advisees want to get the best-informed doxastic attitude possible, and not one which is merely supported by their own evidence. The activity that consists in pursuing this goal is also a form of doxastic deliberation. After all, it is about forming a doxastic attitude. I suggest that we call this kind of doxastic deliberation *public doxastic deliberation*, in contrast to private doxastic deliberation, which is concerned with believing what’s supported by one’s own evidence. It is public in the sense that it goes beyond the realm of the agent’s ‘private’ evidence.

The phrase ‘best-informed doxastic attitude possible’ is vague. In its strictest sense, the most-informed doxastic attitude is one that is supported by the facts, i.e., belief that P where P is true and disbelief that P where P is false. On this picture, advisees pursue the goal of believing what they objectively should

believe P . We find support for this idea in the literature on advice. Judith Jarvis Thomson (1986: 178) writes about moral advice:

“On those rare occasions on which someone conceives the idea of asking for my advice on a moral matter, I do not take my field work to be limited to a study of what he believes is the case: I take it to be incumbent on me to find out what is the case.”

However, as I will argue in section 5.3, this understanding of the epistemic goal of public deliberators doesn’t fit seemingly good practice of epistemic advice. We will discuss a number of cases of epistemic advice in sections 5.4 to 5.7, which will narrow down what exactly it means that public deliberators want the best-informed doxastic attitude possible.

For now, we can give the following account of epistemic advice: Epistemic advisors try to help publicly deliberating agents in their goal to adopt the best-informed doxastic attitude possible, which means, in some cases at least, that this doxastic attitude isn’t supported by the agent’s—i.e., the advisee’s—but by the advisor’s richer evidence.

This definition of epistemic advice seems to be blind to the following kind of case. Imagine you are training to become a statistician and wonder which credence in P is supported by the evidence you have. Even if I know that the evidence you have is actually misleading, I could be said to be helping you when I tell you that you should have credence x to P in light of the evidence you possess. You wanted to know which credence is licensed by the evidence you have since you’re trying to learn how to statistically interpret evidence.⁶

Admittedly, this is a form of epistemic advice, at least in some sense of the term. But it’s not the sense I’m concerned with here. My ‘advice’ in this case is for the sake of helping you to learn how to be epistemically reasonable. This is the same kind of ‘advice’ as (1) is, and I set such cases aside in this paper.

There is one final feature of advice I’d like to point out, which will be relevant in section 5.4. When we give advice to someone in order to benefit them, we do this in a certain way. We tell them that they should do A , so that they follow our advice and do A , which is supposed to be to their advantage. But there are

⁶Thanks to Ralph Wedgwood for making me aware of this kind of case.

other ways to bring someone to do *A*. As Hinchman (2005: 355) points out, we could use fear or flattery to make them do *A*. If I tell you that you should do *A*, while raising my fist in an aggressive way, I'm hardly advising you to do *A*, even if doing *A* would be to your benefit. No, when I give you advice, I rather try to make you do what will benefit you by trying to give you a reason to do *A*. I can do this by explicitly pointing out a reason for you to do *A*. Imagine I'm a doctor and I know that your blood pressure is high, but you don't know this. I tell you:

(2) Your blood pressure is too high. You should exercise.

By telling you that you have high blood pressure, I am pointing out a reason for you to exercise. Furthermore, we can also use should-sentences on their own to give someone a reason. Imagine the same scenario, but I just say:

(3) You should exercise.

Here, I don't explicitly refer to a reason, as I do in (2). However, the act of me uttering (3) seems to be a reason on its own for you to exercise. Hinchman (2005) describes this reason as resulting from the fact that advising is presumptuous. In advising others, we signal that we know better (or at least as well) what needs to be done. Advising is therefore, or so (Hinchman, 2005: 259) argues, an invitation to trust the advisor's judgement on what needs to be done. If I present myself as a trustworthy source on whether to exercise, I give you a reason to exercise by uttering (3).

To conclude, epistemic advice is the attempt to help someone who is deliberating about whether a proposition *P* holds by trying to give them a reason to do what helps them achieve the epistemic goal of adopting the best-informed doxastic attitude to *P* possible.

5.3 Objectivism

With a characterization of epistemic advice on the table, I will in the next sections go through different possible accounts of the truth-conditions of epistemic-advice sentences. For this purpose, I will again use the semantic model of the information-sensitivity of the doxastic 'should' introduced in chapter 2:

Doxastic Should For any context of utterance c that provides a doxastic ordering source, ‘ S should have doxastic attitude D to P ’ is true at c iff having D to P reflects accurately how much i supports P , where i is the information-state provided by c .

The question of what the truth-conditions of epistemic-advice sentences are can be re-described as follows: what kind of information-state is provided by contexts in which doxastic-should sentences are uttered for the purpose of advice? I will call this information-state the *advisory information-state*. In order to figure out what the advisory information-state is, I will again rely on the methodology used in previous chapters, i.e., I will both consult linguistic intuitions and metasemantic considerations.

In chapter 3, I have discussed the objective and the subjective senses of the doxastic ‘should’. In this section, I show that the advisory information-state cannot be the fully realistic information-state, and in the next section 5.4 that it cannot be the subject’s information-state. In sections 5.5 and 5.6, I consider two other information-states, which we have so far not (properly) discussed, and will explain why they aren’t the advisory information-state either. I present my preferred account in section 5.7.

According to the first proposal I want to consider, **Objectivism**, the advisory information-state is the fully realistic information-state:

Objectivism. For any context of epistemic advice c_A , ‘ S should have D to P ’ is true at c_A iff adopting D to P accurately reflects how much the fully realistic information-state supports P .

A context of epistemic advice is a context in which a doxastic-should sentence is uttered for the purpose of epistemic advice. As the reader will remember, the fully realistic information-state is the information-state that, if the doxastic ‘should’ is relativized to it, makes the doxastic ‘should’ carry its objective sense. According to **Objectivism**, advice of the form ‘You should have D to P ’ is thus true iff one objectively should have D to P .⁷

⁷As we have seen in sections 3.4.2 and 3.5, it is probably not the fully realistic information-state, but a slightly different one—such as the truly believable realistic information-state—that the objective doxastic ‘should’ is relativized to. I will ignore this subtlety here.

In order to test **Objectivism** and other accounts of the advisory information-state, I will partly rely on intuitions on variations of the following case. **Police-CCTV-1**, which I introduced in chapter 3, is a slightly altered version of it:

Police. Detective Lester Freamon is pondering about who killed Bob. His evidence is conflicting: it might either have been Omar or Marlo, but it must have been one of them. He looks at his colleague Kima Greggs, who has just entered the room. Unbeknownst to Lester, Kima has solid evidence that it was Omar: her reliable informant Bubbles told her that Marlo was at the other end of town at the time of the murder.

(4) Lester: I really don't know what I should believe anymore.

(5) Kima: You should believe that it was Omar.

Our intuition is that (5) rings true. It is unclear whether **Objectivism** predicts that (5) is true. While the case is set up such that on basis of Kima's evidence it is very likely that it was Omar, it is not stipulated that this is true. After all, even though her informant Bubbles is generally very reliable, he could be wrong in this particular case and actually not have seen Marlo at the other end of town. I think it is telling that we seem to get the intuition that (5) is true, even though the case description leaves it open whether the condition for (5) being true in the objective sense is satisfied. This is a first point telling against **Objectivism**.

A second problem with **Objectivism** is that it seems perfectly fine on some occasions to give advice to suspend judgement or to adopt a credence that neither constitutes belief nor disbelief. This can't be true if the advisory information-state is the fully realistic one. If P is true, one objectively should believe P . If P is false, one objectively should disbelieve P . Since P is either true or false, one always objectively should believe or disbelieve P , but never suspend judgement or adopt a credence that does not constitute belief or disbelief (McHugh 2012: 15). Consequently, on **Objectivism** it is never correct to advise someone to adopt doxastic attitudes other than belief or disbelief.

Third, if the advisory information-state were the fully realistic one, it would be trivial to advise someone not to suspend judgement or not to adopt a credence that doesn't constitute belief or disbelief. According to **Objectivism**, advice of

the form ‘You should not suspend judgement on P ’ or ‘You should not have a credence in P that is not equivalent to believing or disbelieving P ’ would always be true and thus uninteresting. Contrary to this, it does not always seem to be trivial. Take the following case introduced by Turri (2012):

Withholding mathematicians (WM). One hundred eminent mathematicians discuss, outside of your earshot, whether a certain axiom entails a certain claim. Let’s call the proposition that the axiom entails the claim ‘ M ’. After their discussion, they all tell you that you shouldn’t withhold on M . You have no other evidence regarding M .⁸

When the mathematicians tell you that you should not suspend judgement on M , this carries interesting information. It implies that the mathematicians think that they have evidence that supports believing or disbelieving M . If they didn’t think so, they would tell you to suspend judgement on M . However, if the ‘should’ were objective, it wouldn’t carry this implication. On this reading, anyone, independent of what they believe about their evidence, is in a position to know and therefore in a position to reasonably assert what the mathematicians are asserting.

Fourth, metasemantic considerations also speak against **Objectivism**. As I pointed out in chapter 2, which information-state is provided by the context in which a should-sentence is uttered is either determined by the speaker’s intentions, the utterance’s pragmatic purpose, or by what’s conversationally salient. At first, it seems to be that the pragmatic purpose of epistemic advice speaks in favour of **Objectivism**. As I mentioned above, on a prima facie plausible interpretation, ‘the best-informed doxastic attitude possible’ is the doxastic attitude that is supported by the fully realistic information-state. This is the doxastic attitude an omniscient agent has and one couldn’t be better informed than someone who is omniscient. On this reading, the epistemic goal of advisees is to believe all truths, or at least hold true beliefs about whether P for every proposition P they are considering. Since advisors want to help their advisees with their epistemic goals, this would thus imply that advisors are talking about what advisees should objectively believe.

⁸This is, almost verbatim, Comesaña’s (2013) description of the case. I have only changed the propositional letter. I will discuss this case in more detail in section 5.8.

However, this view of the goal of epistemic advisees must be wrong. Otherwise, it would never be considered helpful when one is advised to adopt another doxastic attitude than belief or disbelief, and this is clearly absurd. Therefore, the goal of the advisee must be to have a doxastic attitude which is licensed by an information-state that is different from the fully realistic one and that can license doxastic attitudes other than belief. If the pragmatic purpose of epistemic advisors in such contexts is not to help advisees believe all truths, reasonable advisors are also not *intending* to speak relative to the fully realistic information-states in such contexts. These two contextual features suggest that **Objectivism** is false. Finally, I don't see a reason why we should assume that the fully realistic information-state is contextually salient; if anything, the other two features make another information-state salient.

For these four reasons, I reject **Objectivism**. Before moving on to the next proposal, I would like to address an issue about the nature of the advisory information-state. I have just argued that if the pragmatic purpose of advice cannot be to adopt a doxastic attitude licensed by the fully realistic information-state, then a reasonable advisor won't intend to speak relative to it. This raises the questions what happens when an *unreasonable* advisor intends to speak to it or another information-state that does not suit the pragmatic purpose of advice. If it is speaker intentions that determine the information-state, then in such situations the 'advisor' would say something that is true, but not something that would serve the purpose of their utterance qua epistemic advice.

In order to avoid going into unnecessary details, I won't settle here whether it is speaker intentions or pragmatic purpose that determine the information-state. Rather, I make the following conditional claim: if the intentions-based account is correct, then what I am interested in in this chapter is actually not which information-state is provided by contexts of epistemic advice, but rather which is provided by contexts of *good* epistemic advice, where the advisor's intentions are reasonable with respect to the pragmatic purpose of their utterance. In other words, what I am looking for is not which information-state all epistemic advisors are always talking to, but which they should be talking to as epistemic advisors. The reason is that if the intention-based account is right, then there won't be anything like a uniform advisory information-state and a uniform semantics for epistemic-advice sentences, since which information-state is provided by a context

of epistemic advice would simply be whatever information-state the respective advisor intends to talk to, which could be any. I won't make this conditional claim explicit again in the remainder of the chapter, but readers should keep it in mind. For this reason, I will also only consider reasonable, but not unreasonable, speaker intentions as a metasemantic factor determining the advisory information-state.

5.4 Subjectivism

According to the subjectivist, the advisory information-state is the subject's body of evidence:

Subjectivism. For any context of epistemic advice c_A , ' S should have D to P ' is true at c_A iff adopting D to P accurately reflects how much S 's information-state supports P .

In other words, according to **Subjectivism** advice of the form ' S should have D to P ' is true iff S should subjectively have D to P . As I explained in chapter 3, the subjective 'should' can be characterized in different ways, depending on what we take to be S 's information-state. It can be such that one should subjectively believe P iff one is justified to believe P , iff one's evidence supports believing P , iff one's belief that P is the product of a reliable process, or iff one would know P if one came to believe P . I think that on none of these readings of the subjective doxastic 'should' **Subjectivism** turns out right.

To see this, let's think about **Police** again. Is (5) true according to **Subjectivism**? If we specify the case such that Lester neither would come to know nor is justified to believe that it was Omar, it predicts that (5) is false:

Police-New. Everything is as in **Police**. Furthermore, Lester is new in the police department and one of his colleagues has just told him, wrongly, that Kima is often wrong in judging who committed a crime.

(5) Kima: You should believe that it was Omar.⁹

I submit that (5) appears to be as true in **Police-New** as in **Police**, but that **Subjectivism** predicts that it is false. Lester's relevant evidence in **Police-New**

⁹Kolodny and MacFarlane (2010: 119) discuss an analogous variation of **Miners**.

is that Kima told him that it was Omar and that another colleague told him that Kima is not reliable when judging who committed a crime. This evidence hardly supports believing that it was Omar. Furthermore, that Lester has information that Kima is not to be trusted when it comes to judging who killed Bob works in general as a defeater on Lester being justified to believe that it was Omar. This should be so on any plausible account of justification, including sufficiently sophisticated reliabilist ones. For example, on Goldman's (1986: 111) reliabilism, it is a necessary condition for a belief to be justified that the subject does not believe that the process that produced his or her belief is unreliable. Finally, since justification is necessary for knowledge, Lester would also not come to know that it was Omar if he believed it. To conclude, under no reasonable reading of the subjective 'should' comes (5) out true.¹⁰

Notice that Lester's information is incorrect though. Kima usually interprets the evidence correctly and, in this case as well, has come to the conclusion supported by her evidence. It doesn't seem as if adding the detail that Lester has received the mistaken information about her reliability makes (5) false. You might not share this intuition because you might think that the (misleading) evidence that Kima is unreliable works as a defeater for it to be the case that Lester should believe that it was Omar. If this is so, you're stuck with thinking about what Lester should *subjectively* believe. The **WM** case I described in the previous section might be better suited than **Police-New** to convince you that advisors are not talking about what their advisees subjectively should believe. It is structurally alike to **Police-New** in that the mathematicians apparently correctly recommend a doxastic attitude to you, the advisee, that is not supported by your evidence.

In any case, I think we can further back up the case against **Subjectivism** with metasemantic considerations. If the purpose of uttering (5) is to epistemically advise Lester, the goal of this speech act is that Lester adopts the best-informed doxastic attitude to <Omar killed Bob> possible. If the advisor's evidence is more informative than the advisee's, then the epistemic goal of the advisee cannot be to adopt a doxastic attitude supported by their own, but not the advisor's

¹⁰This is premature. As I have pointed out in section 2.7, on one possible reading of the subjective 'should', it is relativized to the available evidence. I will address this possible reading separately in the next section 5.5.

evidence, since this doxastic attitude wouldn't be the best-informed possible. In order to serve the pragmatic purpose of her utterance, which is to help Lester with his epistemic goal, Kima should thus not speak relative to Lester's information-state. If Kima is reasonable, she's thus not intending to speak relative to Lester's evidence when uttering (5). Finally, given that both the pragmatic purpose of Kima's utterance and her speaker-intention are not linked to Lester's information-state, there is no reason to assume that the latter is conversationally salient. To conclude, **Subjectivism** is also wrong in light of metasemantic considerations.

There is an oddity of **Police-New**, which might undermine it as a counterexample to **Subjectivism**. As mentioned in section 5.2, advisors want to benefit their advisees by trying to give them reasons. Can we see Kima as trying to give a reason to Lester? As the case is set up, Lester does not have all-things-considered reasons to believe that it was Omar after hearing (5). So, it seems like Kima has failed to give him a reason.

However, it is important that we distinguish between all-things-considered and *pro tanto* reasons here. Hinchman (2005) himself, from whom I have taken this condition on advice, writes that one advises to do *A* only if one "intends that [the advisee] gain[s] access to a *pro tanto* reason to" do *A* (ibid.: 361). Now, Kima telling Lester that he should believe that it was Omar is arguably giving Lester a *pro tanto* reason to believe that it was Omar, but a reason that is outweighed, or rather defeated by, the misleading information Lester has about Kima.

Could one argue for a stronger condition on (epistemic) advice, according to which in order to give advice to do *A* to an advisee, one needs to intend to give them an all-things-considered reason to do *A*? I think this condition is unreasonably strong. I admit though that what seems plausible about this suggestion is the following consideration: if the goal of advice is to benefit one's advisee by making them do what one advises them to do, then it seems insufficient to intend to give them just a *pro tanto* reason. If advisees are reasonable, they won't do what one advises them to do as long as they don't have an all-things-considered reason to do it. Now, this consideration, however, only supports the weaker condition that advising to do *A* requires intending to give the advisee a reason such that *in combination with their other reasons*, the advisee ends up with an all-things-considered reason to do *A*. If this last condition is satisfied, even giving a reason to do *A* that falls short of being an all-things-considered

reason to do A can achieve the goal of advice to make a (reasonable) advisee do A . Let's call this the *ATC condition* and look at it in more detail.

First, notice that the ATC condition only disqualifies one version of **Police-New** from being epistemic advice. We have to distinguish between a version in which Kima knows that Lester was misinformed about how reliable she is and a version where she doesn't know this. In the first version, she can't reasonably utter (5) with the intention that Lester ends up with an all-things-considered reason to believe that it was Omar. Kima knows that Lester has other reasons, which in combination with the reason she's giving him won't make him have an all-things-considered reason to believe that it was Omar. However, in the second version, where she's ignorant about Lester's misinformation, she can reasonably utter (5) with this intention. She can reason that if an experienced colleague like her tells Lester that he should believe a proposition P , then this makes him, *ceteris paribus*, have an all-things-considered reason to believe P . Admittedly, Kima's intention won't be realized; but the utterance of (5) will still be advice, since what matters is the intention, not whether it is realized. Demanding that an utterance actually has to bring it about that the advisee has an all-things-considered reason to do A in order to be advice to do A is too much. Otherwise, a report like the following one would be inconsistent:

- (6) He advised me to try the peanut butter sandwich because it is very tasty. But I'm allergic to peanuts, so I didn't eat it.

If advice required succeeding in giving the advisee an all-things-considered reason, then the first sentence and the second couldn't possibly both be true. But they seem to be both true.

Second, I think the ATC condition doesn't apply to advice, at least to epistemic advice. To see this, consider a case where I advise you to have a credence of x in P . On the ATC condition, this can only be advice if I intend to thereby bring it about that you have an all-things-considered reason to adopt credence x in P . I contend that I will intend this only in very rare cases. I can only (reasonably) intend this if I believe that you will then have an all-things-considered reason to have credence x in P . That is, I need to believe that the following rational constraint holds on your probability function Cr :

Full Trust $Cr(P|I've\ advised\ you\ to\ adopt\ credence\ x\ to\ P) = x.$

A scenario where this is given is when you take me to be perfectly reliable. However, usually you will assume that there is a chance that I'm wrong, however small that might be. As a result, your conditional credence in P given that I've told you that you should have credence x to P will in such cases usually be somewhere between your unconditional credence in P and x . The more reliable you think I am, the closer to x it will be. True, there are cases where **Full Trust** is satisfied even though you don't think I'm generally perfectly reliable, for example when your unconditional probability to P happens to be x anyway. But in all other cases where you don't expect me to be generally perfectly reliable, you won't have all-things-considered reason to adopt credence x to P , but rather to adopt a credence between x and your unconditional credence in P . The ATC condition entails that in all such cases I can't, reasonably, give you epistemic advice, since I can't, reasonably, intend to put you in a position to have an all-things-considered reason to adopt credence x to P . I find it highly implausible that reasonable epistemic advice to adopt a certain credence only exists in such rare cases.

I take it that we've established that Kima's utterance of (5) is a form of epistemic advice, at least in one of the two versions of **Police-New**, and very likely in both. Given this, we can back up the intuitive case against **Subjectivism** with the metasemantic considerations that I mentioned above, which derive from the assumption that Kima is epistemically advising Lester.

5.5 Available Evidence

This position is a cousin of **Subjectivism**. It states:

Available Evidence. For any context of epistemic advice c_A , ' S should have D to P ' is true at c_A iff adopting D to P accurately reflects how much the evidence available to S supports P .

Dowell (2013: 160) supports the idea that, in contexts where an agent is in practical deliberation about what they should do, this 'should' is relativized to the evidence *available* to the agent, and not simply the evidence which is in their

possession. If we assume that the subjective ‘should’ is relativized to the subject’s available evidence—as when we join **Justification** with Gibbons’ (2013: 39) quite externalist account of justification¹¹— **Available Evidence** is simply another way to specify **Subjectivism**.

Available Evidence is attractive for our cases dealing with epistemic rather than practical normativity since it predicts that (5) is true in at least some specifications of **Police-New**. Under normal circumstances, the evidence available to Lester includes Kima’s. He just needs to ask her why she thinks that he should believe that it was Omar, in order to access the information that Bubbles told her that Marlo was at the other end of town. The total relevant evidence available to Lester under normal circumstances in **Police-New** is:

- *E1*: It was either Omar or Marlo.
- *E2*: Kima told him that he should believe that it was Omar.
- *E3*: His colleague told him that Kima is unreliable when it comes to judging who committed a crime.
- *E4*: The generally very reliable informant Bubbles told Kima that Marlo was at the other end of town at the time of the murder.

While *E3* undermines *E2*, they are both irrelevant in light of *E1* and *E4*. Here’s one way to put this: the probability of Omar being the murder conditional on *E1* and *E4* is the same conditional on *E1-E4*. Since *E1* and *E4* strongly support that it was Omar, the total evidence available to Lester thus supports it. Accordingly, (5) is true on **Available Evidence** under normal circumstances in **Police-New**.

But circumstances aren’t always normal. One specification of **Police-New** looks as follows:

Police-Disappears. Things are as in **Police-New** (p. 130). Furthermore, Kima disappears right after uttering (5).

To my ears, (5) still rings true in **Police-Disappears**. However, (5) is not true on **Available Evidence**, since the fact that Bubbles told Kima that it cannot

¹¹For more on this, see subsection 3.4.1.

have been Marlo is not evidence that is available to Lester. Kima disappears, so there's no way for him to access this evidence.

I'm not sure that everyone shares my intuition that (5) is true in **Police-Disappears**. However, I think that the same metasemantic consideration that counts against **Subjectivism** undermines **Available Evidence**. Doxastic deliberators want to tap into their advisors' knowledge. In light of this, it would be absurd if advisors, who tries to help their advisees with their deliberative goals, would limit themselves to the evidence that is available to their advisees. By drawing on those parts of their—the advisors'—evidence that aren't available to their advisees, advisors can recommend advisees a better informed doxastic attitude. The pragmatic purpose of advice thus requires advisors to speak to an information-state that is *not* limited to evidence available to the advisee. Accordingly, reasonable advisors won't be intending to speak relative to such an information-state. As in previous cases, these two factors indicate that such an information-state is also not conversationally salient.

5.6 Advisor's Evidence

The discussion in the previous section suggests the following view:

Advisor's Evidence. For any context of epistemic advice c_A , ' S should have D to P ' is true at c_A iff adopting D to P accurately reflects how much the advisor's evidence supports P .

This gets us the right results for all the discussed cases so far since Kima's evidence supports believing that it was Omar, meaning that (5) turns out true on **Advisor's Evidence**. It also fits with the theoretical consideration that it is absurd, given the aim of advice, for advisors not to look at their own evidence when giving advice.¹²

¹²The idea that the practical 'should' is in some contexts sensitive to the speaker's, i.e., the advisor's, evidence has been discussed in the literature. See Kolodny and MacFarlane (2010), Björnsson and Finlay (2010), Dowell (2012, 2013).

However, **Advisor's Evidence** faces a problem. Consider the following case:

Police-Reverse. Lester and Kima both know that it was either Omar or Marlo who killed Bob. Kima was told by her generally very reliable informant Bubbles that Marlo was at the other end of town at the time of the murder. Lester knows that Bubbles told this to Kima, but Lester also knows that Bubbles is mistaken to think that he has seen Marlo at the other end of town. Lester has not shared this information with Kima. Kima walks into the room, where Lester is pondering about the case.

(5) Kima: You should believe that it was Omar.

(5) seems to be false here. One of the reasons is that following it would make Lester adopt a belief that it is not supported by his own, *better* evidence. A parallel example from practical advice should make this even clearer. Imagine you want to go from St Andrews to Edinburgh by the cheapest mode of transport. You know that this usually would be by bus, but that the bus company doubled the price today because of the Edinburgh Festival. Not knowing about this, I tell you: 'You should take the bus'. If advice was correct if it was true in light of the advisor's evidence, my utterance would be true. Contrary to this, my advice appears to be clearly false.

Furthermore, it's certainly not the goal of advisees to have a doxastic attitude that is supported by their advisor's evidence, but not supported by their own information-state, where this information-state is better. It follows that Kima's information-state in **Police-Reverse** is not the one Kima should be talking to in order to serve the pragmatic purpose of her utterance. However, since she is apparently not aware that Lester has better evidence than her, she seems to reasonably intend to speak relative to her own evidence. **Police-Reverse** could therefore be the first case we're encountering in which the information-state suggested by the pragmatic purpose and the one the advisor reasonably intends to speak to come apart. If it is the speaker's reasonable intention that determines which is the contextually provided information-state and Kima really intends to speak to her own information-state, then (5) in **Police-Reverse** is actually true.

I think we can save our intuition that (5) is false in **Police-Reverse** and avoid

having to decide between reasonable speaker intentions and pragmatic purpose as determining the information-state. As I argue in the next subsection, the contextually provided information-state is the collective information-state, the information-state that pools the advisor's and the advisee's evidence. On this account, (5) is false in **Police-Reverse**. Moreover, the account is compatible with Kima's intention determining the context's information-state since Kima might just as well be intending to speak to this collective information-state when uttering (5) in **Police-Reverse**. Since she falsely believes that Lester has no further relevant evidence on whether it was Omar, she falsely thinks that the collective information-state supports believing that it was Omar and therefore utters (5).

Before turning to my preferred account, I'd like to discuss one possible modification of **Advisor's Evidence**. On this version of it, the advisory information-state is the information-state of an advisor who is better informed than the advisee. Kima's information-state in **Police-Reverse** couldn't be the information-state that, according to this account, (5) in **Police-Reverse** is relativized to, since Kima is not better informed than Lester. However, what would the advisory information-state be instead? If there was none, (5) would have no content. But if (5) is false, it must have content. In conclusion, the refined version of **Advisor's Evidence** must also tell us what the advisory information-state is where the advisor is not better informed. Here is such an account:

Better-Informed Evidence. For any context of epistemic advice c_A , ' S should have D to P ' is true at c_A iff adopting D to P accurately reflects how much i supports P , where i is the information-state of the advisor if the advisor is better informed and that of the advisee S if the advisee is better informed.

Better-Informed Evidence gives us the correct result for all the cases discussed

so far. It is not without difficulties either, though. The following case shows this:

Police-Equals. Kima only knows that either Marlo or Omar killed Bob. Lester only knows that either Omar or Avon Barksdale killed Bob. Kima says to Lester:

- (8) Kima: You should suspend judgement on whether it was Omar who killed Bob.

The first problem is that **Better-Informed Evidence** does not determine an advisory information-state for (8). Kima's and Lester's evidence are equally good, so neither information-state is better than the other. We could fix **Better-Informed Evidence** by stipulating that in cases of ties either the advisor's or the advisee's information-state is the advising one. One problem with this is that I'm a bit at loss to see which of the one to choose. Choosing either seems arbitrary.

Second, if we chose either of them, (8) turns out true. Both information-states support suspending judgement about whether it was Omar since on either information-state it could as well have been someone else (Marlo or Avon Barksdale). However, there's good reason to believe that (8) is false. It is bad advice. Kima seems not to be interested in what Lester's evidence concerning the case is. To give advice to someone without first checking whether they might have evidence that could put one in an epistemically better position to give advice is arguably wrong.

Third, **Better-Informed Evidence** requires making precise what makes an information-state 'better' than another one. In **Police-Equals**, it is easy to measure the relative quality of the two information-states. Intuitively, they are equally informative on the issue of who killed Bob. But the more complex, that is, the more like in the real world, the involved information-states get, the harder it will be to make comparisons. I'm not saying that this is impossible. But providing this adds to the complexity of our account of correct epistemic advice. If we could give a simpler model, this would be preferable.

5.7 Collectivism

The following account can deal with all these three issues:

Collectivism. For any context of epistemic advice c_A , ‘ S should have D to P ’ is true at c_A iff adopting D to P accurately reflects how much the collective evidence supports P .¹³

The collective evidence is the result of pooling the advisor’s and the advisee’s evidence. Pooling roughly means combining or putting together the advisor’s and the advisee’s evidence. The effect is such that where one or both individuals have evidence the other doesn’t possess, the collective evidence is better-informed than their respective bodies of evidence. I will discuss different possible ways of pooling evidence in the next chapter.

Collectivism is superior to **Better-Informed Evidence**. First, it avoids an arbitrary choice of one of the information-states in cases where the two are equally good. In every case, the account simply pools them. Second, on **Collectivism**, (8) in **Police-Equals** is false. The doxastic attitude licensed by the intersection of Kima’s and Lester’s evidence is believing that it was Omar since the intersection of the proposition <It was either Omar or Marlo> and the proposition <It was either Omar or Avon Barksdale> is the proposition <It was Omar>. The account also makes sure that in order to give true epistemic advice, the advisor is urged to look at the advisee’s information-state, and not only their own information-state, before giving advice. This includes cases where the advisee’s information-state is actually better than the advisor’s, but also cases such as **Police-Equals** where the information-states are equally good. Third, it avoids the problem of how to measure the quality of information-states, since it combines the two information-states, rather than compares them.

Collectivism also has none of the problems any of the other mentioned accounts have. Unlike **Subjectivism** and **Available Evidence**, it includes the advisor’s evidence into the advisory information-state. Thus, (5) is true

¹³A related account is considered by Kolodny and MacFarlane (2010) and supported by Björnsson and Finlay (2010: 13) for the truth-conditions of practical advice. While Björnsson and Finlay only discuss examples of practical advice, Steve Finlay told me in personal communication that they consider their view as applying to epistemic advice, too. The present paper can be seen as taking up the task of defending their view as it applies to epistemic advice.

in **Police-New** and **Police-Disappears** according to **Collectivism**, which is what we want. Unlike **Objectivism**, it allows for the truth of epistemic-advice sentences that recommend doxastic attitudes other than (dis)belief, since it does not postulate that the advisory information-state is always the fully realistic one. Finally, unlike **Advisor's Evidence**, it doesn't make an epistemic-advice sentence true if it is true in light of the information-state of an advisor that is less-informed than their advisee since **Collectivism** also includes the advisee's evidence into the advisory information-state.

Finally, **Collectivism** is supported by metasemantic considerations. The epistemic goal of deliberating agents is to adopt the best-informed doxastic attitude to P possible. As explained in section 5.3, this cannot be the goal of believing all truths since otherwise advice to adopt doxastic attitudes other than belief or disbelief could never be considered help with respect to deliberating agents' epistemic goals. Furthermore, because advisees want to tap into their advisor's evidence, it is their goal to adopt a doxastic attitude supported by an information-state containing both the evidence that they possess and that possessed by their advisor. Finally, in cases where the advisee's information-state contains further relevant evidence that the advisor's does not, the doxastic attitude supported by the advisor's, but not the advisee's, information-state can hardly be called best-informed. Advisees want to tap into their advisor's evidence, but if they have further relevant evidence, they obviously don't want to ignore this when forming their doxastic attitude to P . We can respect all this if we assume that the epistemic goal of deliberating agents is to adopt the doxastic attitude that is supported by the collective information-state. The pragmatic purpose of epistemic advice therefore requires advisors to speak relative to this information-state, and they accordingly intend to speak to it if they are reasonable. That these two contextual features suggest that the collective evidence is the information-state provided by contexts of epistemic advice provides metasemantic support for **Collectivism**.

5.8 Resolving Turri’s Puzzle about Withholding

I have argued that the collective doxastic ‘should’ is the ‘should’ occurring in epistemic advice.¹⁴ This is philosophically noteworthy because it fills a gap in the literature concerning the truth-conditions of epistemic-advice sentences. Furthermore, understanding what it means to give *correct* epistemic advice also helps us to better understand the practice of giving epistemic advice, which has so far been overlooked in the literature.

Apart from these rather big-picture considerations, I think that **Collectivism** also has a more concrete philosophical cash value. It can solve a puzzle that John Turri describes in his paper ‘A Puzzle about Withholding’, surrounding the case **WM** mentioned above. The puzzle arises from three assumptions, which he takes to be common in epistemology:¹⁵

- *Triad*: For every proposition P that a subject S considers, S can adopt three and only three doxastic attitudes towards P : believe it, disbelieve it, or withhold.¹⁶
- *Optimism*: Given any proposition P and any body of evidence E , E will justify at least one doxastic attitude towards P .
- *Evidential Propriety*¹⁷: The epistemic propriety of a subject’s doxastic attitude is entirely a function of the subject’s evidence.

These three assumptions come into conflict in **WM**:

Withholding mathematicians (WM). One hundred eminent mathematicians discuss, outside of your earshot, whether a certain axiom

¹⁴A longer, self-standing version of the present section has been published under the title ‘Resolving Turri’s Puzzle about Withholding’ in *dialectica* (vol. 70, issue 2).

¹⁵In the following, I will refer to the set of these views as ‘the three assumptions’.

¹⁶One might object to this assumption that it ignores credences, i.e., more fine-grained doxastic attitudes. If one thinks that not every credence is an instance of one of the coarse-grained doxastic attitudes, belief, disbelief or withholding, one will deny *Triad* and claim that besides the three coarse-grained doxastic attitudes, a subject could also adopt some credence to a proposition. However, this objection does not undermine the following discussion. Turri’s puzzle arises in a similar way if one reads the three assumptions only as claims about coarse-grained doxastic attitudes, and not as claims about all doxastic attitudes in general.

¹⁷Turri (2012) actually calls this assumption *Evidentialism*. I have chosen another term to avoid confusion with the different position **Evidentialism**, which I introduced in section 3.2.

entails a certain claim. Let's call the proposition that the axiom entails the claim 'M'. After their discussion, they all tell you that you shouldn't withhold on M. You have no other evidence regarding M.

Turri (2012: 359) states that the following two claims seem true:

- (9) You should not withhold on M in **WM**.
- (10) You should neither believe nor disbelieve M in **WM**.

Turri's argument for (9) is concise: "It stands to reason [...] that if all the Mathematicians say that withholding is not the thing to do, then you shouldn't withhold" (ibid.). His argument for (10) rests on *Evidential Propriety*. Since your evidence neither supports believing M nor disbelieving M, you should neither believe nor disbelieve M.

Turri argues that **WM** creates a puzzle about withholding since it apparently shows the possibility of an epistemic impasse. An epistemic impasse is a situation in which one is not permitted to adopt any doxastic attitude to a proposition P. According to (9), you should not withhold and according to (10), you should neither believe nor disbelieve. From (9) and (10), Turri concludes that

- (11) You are neither permitted to believe or disbelieve M nor permitted to withhold on M in **WM**.¹⁸

Since according to *Triad*, you can only adopt belief, disbelief or withholding towards M, you are thus not permitted to adopt any doxastic attitude to M. Thus, **WM** is an epistemic impasse. This is incompatible with *Optimism*. According to *Optimism*, given the proposition M and your evidence in **WM**, there is at least one doxastic attitude you are justified to adopt to M.

Turri goes through several possibilities of resolving this puzzle by either giving up *Triad*, *Evidential Propriety* or the intuitive verdict that (9) is true.¹⁹ He rejects all of these options and concludes: "None of the proposed solutions to the puzzle seems fully satisfactory. I, for one, am left puzzled" (ibid.).

I won't discuss Turri's arguments against the possible resolutions he considers. My alternative way of resolving the puzzle should be clear: I claim that (9)

¹⁸(11) follows from (9) and (10) if one assumes, as Turri obviously does, that if one should not do A, one is not permitted to do A. I will challenge this assumption further below.

¹⁹He does not consider rejecting *Optimism*.

is true because the ‘should’ in it is collective. When the mathematicians are giving epistemic advice, they talk about what you should believe in light of the collective evidence. I refer to the proposition expressed by (9) when uttered by the mathematicians as (9_C):

(9_C) In light of the collective evidence, you should not withhold on *M*.

The ‘should’ in (10), on the other hand, is subjective. Since (9_C) does not entail that you are not permitted not to withhold, (9) and (10) together do not entail (11).

The reason for this is as follows. That in light of your evidence you should not do *A* entails that you are not permitted to do *A*.²⁰ However, that you should not do *A* in light of a body of evidence that contains evidence that you don’t possess does not entail that you are not permitted to do *A*, as I will argue for shortly. That is, the subjective ‘should’ is tied to permission, but the collective ‘should’ isn’t. In other words, (10) entails (11) in combination with (9_S), but not (9_C):

(9_S) In light of your evidence, you should not withhold on *M*.

(9_S) is false. As Turri remarks himself, your evidence neither supports believing *M* nor $\neg M$. Consequently, you should withhold in light of it. You might think that (9_S) is actually true since after the mathematicians uttered (9), you gathered new evidence supporting that you should not withhold: the experts told you not to withhold. What speaks against this line of reasoning is that on the natural picture I am assuming throughout this thesis, which doxastic attitude one should have to a proposition *P* in light of a body of evidence *E* depends on the evidential support relation between *E* and *P*. That the mathematicians uttered (9) is part of your evidence, but this additional evidence neither supports nor speaks against *M*, as can be seen by the fact that we are at a loss to say whether it supports *M*

²⁰This needs qualification. As explained in chapter 2, ‘should’ is a weak deontic necessity modals in comparison to the strong deontic modal ‘must’. It is agreed that $\langle S \text{ must do } A \rangle$ (in its subjective sense) entails $\langle S \text{ is obliged to do } A \rangle$ and therefore, since obligation is the dual of permission, that $\langle S \text{ must not do } A \rangle$ entails $\langle S \text{ is not permitted to do } A \rangle$. However, ‘should’ is thought to be weaker and not to entail obligation, from which follows that $\langle S \text{ should not do } A \rangle$ might not entail $\langle S \text{ is not permitted to do } A \rangle$. (Harman 1993 is one among many to have made this point.) I ignore this subtlety here. If ‘should’ is not the dual of ‘is permitted’, Turri could simply restate the problem in terms of ‘must’.

or speaks against it. Therefore, even once you have gathered the evidence that the mathematicians uttered (9), you should withhold on M in light of your evidence.

Now, since (9) in **WM** is only true in the sense of (9_C), but not (9_S), **WM** does not constitute an epistemic impasse because (9_C) and (10) do not entail (11). To see that such an entailment relation between the collective ‘should’ and permission does not hold, imagine that you simply formed the belief that M (or $\neg M$; which is irrelevant). It seems fairly clear to me that you are not permitted to do this. You have no evidence supporting M , so you’re arbitrarily choosing a belief on whether M holds. But arbitrarily choosing beliefs is hardly epistemically permissible. If it were, it would be OK for you to simply form the belief that the numbers of stars in the universe is even or that the fair coin you are about to flip will land heads up. This is so, even if someone *else* happens to have excellent evidence that the numbers of stars is in fact even or that the coin will land heads up. Thus, whether you are permitted or not to adopt a doxastic attitude seems to be merely a function of your own evidence, but not of that of others. We have accordingly no reason to assume that the fact that you should not, in the collective sense, withhold on M entails that you are not permitted to withhold on M .

A worry one might have about my solution is whether (9) can really be called advice and whether it is hence plausible to claim that the ‘should’ in it must therefore be collective. As pointed out in section 5.2, it is widely assumed that advice requires that advisors intend to give advisees a pro tanto reason to do whatever they advise them to do. If the mathematicians utter (9), this doesn’t seem to give you any kind of reason not to withhold on M . It’s not like the utterance of (9) gives you a pro tanto reason not to withhold, which is then overridden by another reason. Since it doesn’t change your evidence with respect to M at all, it is hardly a reason for you to stop withholding on M at all.

These considerations do not only make it doubtful whether (9) is a form of epistemic advice, but also raise the question of why the mathematicians utter it in the first place. If they don’t want to advise you, what else are they trying to achieve? Turri doesn’t tell us more about the case than what is said in **WM**, so it is underspecified with respect to the advisors’ intentions. I see two possible reasons for the mathematicians to utter (8) when it carries its collective meaning (9_C). First, as I explained in section 5.3, the mathematicians can use (9) to simply

imply that they have evidence that determines whether M is true. Where making this implication is their sole intention, they do not try to give you a pro tanto reason not to withhold on M . Under this specification of **WM**, (9) would not be epistemic advice, but nonetheless contain the collective doxastic ‘should’.

Second, things in **WM** could also be as follows. Imagine that what the mathematicians are, and only are, discussing outside your earshot are the merits of a recently suggested proof for M and that, so far, M has neither been proved nor disproved. They can therefore only come to two possible conclusions after discussing this proof: that you should (collectively) believe M or that you should keep on withholding on M . They come to conclude that the proof works and that you should believe M . The crucial point is that they reasonably, but mistakenly, believe that you know that they have been discussing this proof. Furthermore, they reasonably, but mistakenly, believe that you have been told by a rival, evil mathematician that they cannot be trusted and that whatever they say, you should maintain withholding on M . In such a situation, the (good) mathematicians can reasonably intend to give you a pro tanto reason not to withhold on M by uttering (9). They assume that on hearing (9) you get a reason to believe that the collective evidence supports M because they think that you know that the only options are that the evidence supports believing or withholding on M and that you therefore will infer from (9) that they believe that the collective evidence supports M . They don’t think that you have an all-things-considered reason not to withhold on M after hearing (9), as they believe that you also have evidence from their rival, evil mathematician that supports withholding on M and weighs against the reason they have provided. To conclude, under this second specification of **WM**, the good mathematicians can reasonably intend to give you a pro tanto reason not to withhold on M by uttering (9), and (9) will therefore be reasonable epistemic advice under these circumstances.²¹

It is important to notice that my view is compatible with *Evidential Propriety*.

²¹You might wonder why I have added an extra complication to this specification by stipulating that you have been misinformed by a rival, evil mathematician. The reason is that otherwise the good mathematicians are giving you an all-things-considered, not just a pro tanto, reason not to withhold by uttering (9). An opponent of my solution to Turri’s puzzle could therefore claim that the (9) in ‘should’ is in fact subjective, and not collective. Don’t get me wrong: I do believe that in such a case the ‘should’ is also collective, since (9) is still advice, but I wanted to show that there is a specification of **WM** where it is uncontroversial that (9) is both epistemic advice and contains a collective ‘should’.

While I claim that which doxastic attitude a subject should adopt, in the collective sense, is partly a function of evidence that is not the subject's, I do not deny that the epistemic propriety of a subject's doxastic attitude is entirely a function of the subject's evidence. I assume that by 'epistemic propriety' Turri means something along the lines of 'justification'; and whether you are justified in adopting a doxastic attitude might very well be the same issue as whether you should adopt it *in light of your evidence*. However, whether you should adopt it in the *collective* sense is another issue and neither entails nor is entailed by whether you are justified to adopt it. Since my view is furthermore compatible with *Triad* and *Optimism*, it thus resolves the puzzle while accepting both (9) and the three assumptions.

There are two more reasons why it matters that justification and epistemic propriety come apart from what one should believe in light of the collective evidence. First, it allows me to acknowledge the main point Comesaña (2013) makes in his solution to Turri's puzzle, while maintaining that my solution is superior to his. Comesaña argues that (9) is false since you are not justified in believing or disbelieving *M*. He concludes that you are and you are *only* justified to withhold on *M*.

I agree with this last point, but it does not follow that (9) is false, at least if it means (9_C). It is tempting to derive from the fact that you are not justified to believe or disbelieve *M* that you should not believe or disbelieve *M*. This is indeed a plausible principle if 'should' is read subjectively, and we can correctly derive from it that (9_S) is false. However, this principle does not apply where 'should' is collective. Thus, Comesaña has only shown that (9_S), but not that (9_C) is false.

The important difference between Comesaña's view and mine is that Comesaña claims that (9) is wrong. He does so because he only considers a subjective reading of 'should', as his focus on whether you are justified to not withhold on *M* in *WM* indicates. The advantage of my view over Comesaña's is that I can explain why we intuitively judge (9) to be true: we correctly read it as expressing (9_C) when the mathematicians utter it. On behalf of Comesaña, one could suggest the following alternative, error-theoretic explanation of our judgement that (9) is true: we read it as expressing (9_S) and mistakenly judge it to be true since we wrongly assume that once the advisors have told us not to withhold, we shouldn't withhold

relative to our evidence. My solution has the advantage that it does not have to posit that we're making such a relatively simple mistake and is thus more in line with the assumption that speakers are competent. If this advantage came at the cost of having to make implausible claims about the meaning of (9), the error theory might be preferable. However, my arguments in support of **Collectivism** have given us good reason to assume that the 'should' in (9) is collective.

Second, the distinction between epistemic propriety and the collective 'should' also prevents my solution from being subject to the following objection. One might wonder whether the evidence-sensitivity of 'should' does not also imply that 'is permitted' is evidence-sensitive. In this case, 'is permitted' can carry a collective sense, too. This would allow us to infer (11) from (9_C) and (10) if (11) meant:

(11*) In light of the collective evidence, you are not permitted to withhold on *M*, and in light of your evidence, you are not permitted to believe or disbelieve *M*.

Personally, I have difficulties hearing 'is permitted' as being relativizable to collective evidence. As suggested above, if it were, someone who knows that a coin about to be flipped will land heads up could correctly say about a person who does not know this:

(12) She is permitted to believe that the coin will land heads up.

(12) sounds off to me. In any case, even if (11) could mean (11*), this wouldn't make **WM** an epistemic impasse. Remember that what is problematic about an epistemic impasse is that it is a situation where there is no doxastic attitude to a proposition *P* that would be epistemically proper for the subject. However, the truth of (11*) does not entail that there is no epistemically proper doxastic attitude for you in *WM*. Whether one is justified in adopting a doxastic attitude might be identical with whether one should, in the subjective sense, adopt it, but is distinct from whether one should, in the *collective* sense, adopt it. Analogously, the fact that in light of some collective evidence, which contains evidence you do not possess, you are not permitted to withhold on *M* does not mean that it is not epistemically proper or justifiable for you to withhold on *M*.

5.9 Conclusion

I have argued for **Collectivism**, the view that epistemic advice of the form ‘You should adopt doxastic attitude D to proposition P ’ is true iff D to P reflects how much the combined evidence of the advisor and the advisee supports P . This reveals a further instance of the information-sensitivity of the doxastic ‘should’: its collective sense. Since no one has so far put forward an account of the truth-conditions of epistemic-advice sentences, **Collectivism** fills a gap in the literature. It also helps us to better understand the so far under-researched practice of epistemic advice. On top of this, it can solve Turri’s (2012) puzzle about withholding.

Collectivism can be precisified in a number of ways. One issue, for example, is what we understand as the collective evidence. In the next chapter, I will give an account of how I think **Collectivism** should be specified.

Chapter 6

Precisifying Collectivism

6.1 Introduction

In the previous chapter, I have proposed **Collectivism**, the view that when we are giving epistemic advice, we talk about what our advisees should believe in light of the collective evidence. **Collectivism** is a big tent and allows for different precisifications. In this chapter, I discuss ways in which **Collectivism** can be precisified. The purpose of this is to provide a more complete account of **Collectivism**, and hence of correct epistemic advice, than I presented in the previous chapter and to detect open questions that merit further research.

Section 6.2 examines how the advisor's and the advisee's evidence is pooled, i.e., which parts of their evidence go into the collective evidence. Section 6.3 investigates whether (and mostly rejects that) the literature on group justification can be of any help to precisify **Collectivism**. Finally, section 6.4 discusses whether epistemic advice of the form 'You should believe P ' is true only if P is true. I argue that there is no such factivity necessity condition on correct epistemic advice and that our intuition supporting such a condition is mistaken. This last insight does not only help us to get a better understanding of correct epistemic advice, but also defends **Collectivism** because accepting said intuition actually pushes us to a view that is incompatible with **Collectivism**.

6.2 How to Pool Evidence

The first big question is what exactly the collective evidence is. The collective evidence is the result of pooling the evidence of the advisor and the advisee. This is a process that happens not only on the inter-subjective level, i.e., when two individuals' evidence is joined together, but also on the intra-subjective level. When we gather new evidence, we add this new evidence to our old evidence. Thus, the process of pooling the evidence of two distinct individuals should in principle not be more mysterious than the process of an individual gathering new evidence. However, evidence can be pooled in different ways. In this section, I will discuss the main options, arguing against some of them, but won't settle which method of pooling evidence a defendant of **Collectivism** should go for. I consider this a task for further research.

How we understand the process of pooling evidence depends on our account of evidence. As mentioned in chapter 2, on one picture of information-states they are sets of worlds. The result of pooling two such information-states would be the intersection of these two sets. Where one agent has evidence the other agent doesn't possess, pooling accordingly results in an information-state containing less worlds than the two agents' information-states respectively contain. This reflects that the collective evidence is more informative than the agents' respective sets of evidence: it has narrowed down the set of possible worlds that might be the actual world.

This picture runs into problems in cases where a subject S 's evidence contains a proposition Q and the other subject S^* 's contains its negation $\neg Q$. The information-state representing S 's evidence contains only Q -worlds (worlds at which Q is true) and S^* 's information-state only contains $\neg Q$ -worlds. The intersection of these two sets is the empty set.

This problem is avoided on factive accounts of evidence, for example the view that one's evidence consists of all and only the propositions one knows.¹ Two bodies of evidence that only contain true propositions cannot contain inconsistent propositions since the truth is not inconsistent.

However, this does not mean that **Collectivism** is committed to the factivity

¹Tim Williamson (2000: Ch. 9) famously argues that $E=K$ and hence that only true propositions can be evidence. Littlejohn (2011) argues that evidence is factive, but denies $E=K$.

of evidence. The issue of how to merge evidence has recently received attention among those who apply formal methods to social epistemology, originating from the literature on judgement aggregation (List and Pettit 2002, 2004). Jon Williamson (2009, 2015) discusses how to merge evidence where a subject's evidence is what is rationally granted by the subject, which can include false propositions. He does not propose a specific merging operator, but suggests certain constraints on such an operator.

One operator meeting the suggested constraints adds all and only those propositions to the collective evidence that are granted by all group members (Williamson 2015: 57). This seems far too restrictive as it wouldn't add to the collective evidence evidence that, for example, a better-informed advisor has, but a less-informed advisee lacks. This merging operator, therefore, does not represent the pooling of evidence we're interested in.

An alternatively proposed operator adds all and only those propositions to the collective evidence that are not contentious and granted by at least some members of the group (ibid.: 58). A proposition is contentious, roughly, if there is disagreement within the group with respect to whether it is part of the evidence or not. The collective evidence created by this operator includes therefore the evidence that better-informed advisors possess, but excludes those propositions Q such that the evidence of one member of the group contains Q , but that of another member contains $\neg Q$. Let's call this process *neutral pooling*.

I think that neutral pooling is not inclusive enough, and favour the view that the collective evidence should include a true proposition that is part of one group member's evidence even if its negation is part of the other group member's evidence. On this view, collective evidence formed from two bodies of evidence E and E^* contains a proposition Q if and only if one of the following conditions are satisfied:

- Q is part of E and E^*
- Q is only part of E (or E^*) and $\neg Q$ is not part of E^* (or E)
- Q is only part of E (or E^*) and $\neg Q$ is part of E^* (or E), but Q is true.

Let's call this *truth-sided pooling*. To see the different predictions of neutral and truth-sided pooling, let's look at the following case.

Misheard. Bob is about to undergo a very complicated surgery, which often goes wrong. His doctor Dr Mayer has told him that she's very good at this surgery and that she's unusually successful in performing it, but Bob has heard from a colleague of Dr Mayer that she sometimes exaggerates her medical track-record quite a bit. Furthermore, Bob has overheard one nurse saying to another: 'Out of 33 surgeries, Dr Mayer was successful 13 times.' However, Bob misheard. His partner Ann, who also eavesdropped on the nurse, heard what he had actually said: 'Out of 33 surgeries, Dr Mayer was successful 30 times.'

When Ann visits Bob the next time, the issue of how risky the surgery is comes up:

- (1) Ann: But Dr Mayer said she is very good at this.
- (2) Bob: Oh, I don't know whether I should believe her.
- (3) Ann: You should believe her.

Mistaken perception is a plausible candidate for false evidence, so on a non-factive account of evidence the proposition <The nurse said that 13 of Dr Mayer's 33 surgeries were successful> is probably part of Bob's evidence. (Let's call this proposition $13/33$ in the following). The proposition <The nurse said that 30 of Dr Mayer's 33 surgeries were successful> is part of Ann's evidence. (Let's call this proposition $30/33$ in the following).

If Ann's and Bob's collective evidence is the result of neutral pooling, then (3) is accordingly false. Ann's and Bob's evidence concerning what the nurse said is inconsistent and therefore both pieces of evidence are removed from the collective evidence. The remaining evidence is that Dr Mayer said she's very successful at this surgery and that one of her colleagues said that she sometimes exaggerates her medical track-record significantly. This collective evidence arguably does not support believing that Dr Mayer is very good at this surgery, but rather to suspend judgement about it. On the other hand, if the collective evidence is the result of truth-sided pooling, (3) is true since the collective evidence then contains the true proposition $30/33$, but not the false proposition $13/33$.

Do we judge (3) to be true in this scenario? What speaks for (3)'s *falsehood* is that advisors should first check what their advisee's evidence is before giving advice

(with the exception of cases where they know that their evidence is conclusive). If Ann had asked Bob and found out that he thinks that $13/33$, Ann should have investigated whether her or Bob's evidence is accurate. Thus, it seems premature of Ann to utter (3) if there is evidence that Bob possesses that contradicts her own evidence.

However, it is consistent with Ann's utterance of (3) being true that it is premature or irresponsible. We can hit truth by luck. In fact, I tend to believe that this is what's going on here. Advisors want to help advisees in their goal of adopting the best-informed doxastic attitude possible. Hence, Ann is certainly intending to speak relative to her information that $30/33$. Considerations concerning speaker intentions therefore suggest that this proposition is part of the collective evidence (and hence that $13/33$ isn't since a body of evidence can't contain inconsistent propositions). For the same reasons, the pragmatic purpose of her utterance to advise Bob, i.e., to assist him with his goal of adopting the best-informed doxastic attitude possible, is best served if the collective evidence contains this crucial information. Thus, (3) is true. The reason that uttering (3) is premature is that Ann didn't check whether Bob doesn't have any true evidence that contradicts her own, which she needed to do in order to *know* that (3) is true.

Things get even more complicated if information-states are not sets of worlds, but a probability distribution over a set of propositions.² Now, for any proposition that both parties' information-states assign a probability, the question arises which probability this proposition is assigned by the information-state representing the collective evidence. For example, if S 's information-state determines that $Pr_S(P) = 0.2$ and S^* 's information-state determines that $Pr_{S^*}(P) = 0.8$, what probability should the collective evidence represented by Pr_C assign to P ?

One natural idea would be to split the difference between the two, such that $Pr_C(P) = 0.5$. However, similar reasoning applies here as when transforming inconsistent evidence into a collective body of evidence. Ann's information-state, if modelled as a probability function, will assign a high probability to $30/33$, while Bob's information-state will assign it a low probability. The idea that Ann intends to speak relative to her (indeed) better evidence would have to be accounted for by the collective evidence being represented by a probability function that

²For this suggestion, see Yalcin (2011: 299), Wedgwood (forthcoming: sec. 3), and Carr (2015: 704).

assigns the proposition in question a high probability. A pooling operator with the following property would respect this: for any proposition P and two subjects S and S^* , where $Pr_S(P) \neq Pr_{S^*}(P)$, $Pr_C(P)$ equals that of the two probability assignments $Pr_S(P)$ and $Pr_{S^*}(P)$ that is the result of a more *veridical* credence revision.

What I mean by ‘veridical’ is best explained within the model of Bayesianism. According to Bayesianism, credence revision consists in two steps: an initial revision of one’s credence in a proposition E caused through some external stimulus, e.g., perception, and then a revision of one’s credences in other propositions through conditionalizing on E (Weisberg 2011: ch. 3.3). On Jeffrey Bayesianism, the credence one assigns to E can be short of 1, as we are assuming if we model bodies of evidence as probability distributions that assign probabilities to propositions other than 1 or 0 (Jeffrey 1983: ch. 11).

Two factors determine how veridical a revision of one’s credence in P is. First, it is the more veridical, the more the initial revision that triggered it was hooked to reality. For example, increasing one’s credence in E on basis of having perceived that E is more veridical than doing so on basis of hallucinating that E . Second, it is the more veridical, the more the conditional probability $Pr(P|E)$ used for strict conditionalization (or the corresponding formula for Jeffrey conditionalization) represents something like the objective probability of P on E . For example, a probability function that assigns a low probability to the conditional proposition $\langle \text{I will win the lottery} | \text{I bought a lottery ticket} \rangle$ is more veridical than one that assigns it a high probability.

These ideas for how to pool probabilistic information-states are only sketches of a full account. Furthermore, the arguments I present in favour of truth-sided or veridical pooling might turn out to fail and a more neutral approach that doesn’t judge whose evidence is better might be the way to go for pooling evidence. Working out a more definitive account of **Collectivism** is an interesting avenue for further research.

6.3 Group Justification and Group Knowledge

As we have seen in the previous section, **Collectivism** leaves room for precisification since we can have different accounts of what the collective evidence is.

Another way to precisify **Collectivism** is to approach the issue through the lense of accounts of group knowledge or group justification. The key suggestion would be the following one:

Group-Justification Collectivism. For any context of epistemic advice c_A , ‘ S should have D to P ’ is true at c_A iff the group is justified to have D to P .

A knowledge-version of this would state:

Group-Knowledge Collectivism. For any context of epistemic advice c_A , ‘ S should believe P ’ is true at c_A iff the group would come to know P if it believed P .

The idea is to recruit the rich literature on group justification and knowledge to give a more specific account of **Collectivism**. The suggestion is natural since **Collectivism** states that advisees should believe what the collective evidence supports and it is plausible to assume that this is what the group is justified to believe or knows. In this section, I will argue that the current accounts of knowledge and group justification are not very suitable candidates for this purpose. However, I show that future research on **Collectivism** can nonetheless benefit from the debate about group justification, in particular Lackey’s (ms) recent contribution.

One decision point for someone who wants to make use of this literature for our present purposes is to choose between **Group-Justification** and **Group-Knowledge Collectivism**. In the next section, I argue against the idea that ‘You should believe P ’ as a piece of epistemic advice is true only if P is true. This would obviously count against **Group-Knowledge Collectivism** since knowledge is factive. For this reason, I will focus here on **Group-Justification Collectivism**

Goldman and Blanchard (2015: sec. 4.3) notice that so far there aren’t many detailed accounts of group justification in the social epistemology literature. In the following, I distinguish three different proposals. First, there is what Lackey (ms) calls *Summativism*. This is the view that group justification is solely determined by the justified beliefs of its members. The most detailed instance of this view

is, according to Lackey (ms), Goldman’s (2014) social-process reliabilism, which states the following:

Social Process Reliabilism. “If a group belief in P is aggregated based on a profile of member attitudes toward P , then (ceteris paribus) the greater the proportion of members who justifiedly believe P and the smaller the proportion of members who justifiedly reject P , the greater the group’s level, or grade, of justifiedness in believing P ” (Goldman 2014: 28).

One first issue with this is that it takes *doxastic*, rather than propositional, justification of group members as input to calculate the group’s justification.³ This would make what an advisee should believe dependent on what they and the advisor believe, and not just what the two are propositionally justified to believe. This doesn’t fit with **Collectivism** as the latter refers to the group members’ evidence, so to what they are justified to believe, and not just what they justifiably believe for the right reasons. The latter is a subset of the former, so a view that makes what an advisee should collectively believe only dependent on the group members’ doxastic justification leaves out important information, in the eyes of **Collectivism**. Just think about a scenario in which evidence supporting P stares the advisor in the face, i.e., they are propositionally justified to believe P , but they don’t happen to actually believe P . If we plug Goldman’s definition of group justification into **Group-Justification Collectivism**, we would end up with the implausible claim that the advisor’s evidence for P is irrelevant to whether their advisee should collectively believe P .

We could fix **Social Process Reliabilism** for our purposes by substituting ‘who have propositional justification to believe (reject) P ’ for ‘who justifiedly believe (reject) P ’ in its formulation. However, this rephrasing is not capable of alleviating the following problem: the principle is too ‘neutral’ in the sense that it gives each member’s propositional justification equal weight. It would predict that Ann’s utterance of (3) is false since it advocates that we have to aggregate

³One reason that this and other accounts of group justification focus on doxastic rather than propositional justification might be that the social epistemology literature has in the past mostly focused on providing an account of group belief (Goldman and Blanchard 2015: sec. 4.1-4.2). From this point, asking what additional property such a belief needs in order to be justified might be the more natural next step.

Bob's justification to reject $30/33$ and Ann's justification to believe $30/33$. This aggregation yields that the group consisting of Ann and Bob is neither justified to believe nor justified to reject this proposition.

A way to respond to this is to point out that, since Goldman's notion of justification is reliabilist, Ann and Bob aren't actually equally justified: Ann's belief is the product of perception, a reliable belief-forming process, whereas Bob's belief is the product of hallucination. Thus, Ann's belief is more justified and the group is therefore justified to believe the proposition in question. The problem is that reliabilism is an account of doxastic, and not propositional justification, which gets us back into the problem with **Social Process Reliabilism** we started with.

A second family of accounts of group justification falls under Lackey's (ms) label *Non-Summativism*. Such views see group justification as detached from the justification of the group's members. An example is Schmitt's (1994: 265) following view:

Joint Acceptance. "A group G is justified in believing $[P]$ just in case G has good reason to believe, where G has reason $[R]$ to believe $[P]$ just in case all members of G would properly express openly a willingness to accept $[R]$ jointly as the group's reason to believe $[P]$."

Against this, Lackey (ms) argues that openly accepting reasons for believing P cannot be a necessary condition on group justification for believing P . She gives the example of the board members of Philip Morris who all have good evidence and thus justification for believing that smoking is unhealthy, but who aren't willing to publicly accept the relevant evidence. **Joint Acceptance** predicts that Philip Morris is therefore not justified in believing that smoking is unhealthy, which is absurd. For the same reasons, **Joint Acceptance** is not a good candidate for specifying **Collectivism**. Whether an advisor's advice to believe a proposition is true should not depend on non-epistemic issues like the group members' practical reasons for and against accepting certain evidence in public.

Finally, there is Lackey's (ms) own account:

Group Epistemic Agent. A group, G , justifiedly believes that P iff

- (i) a significant percentage of the operative members of G (a) justifiedly believe

that P , and (b) are such that adding together the bases of their justified beliefs that P yields a belief set that is coherent, and

- (ii) full disclosure of the evidence relevant to the proposition that P , accompanied by rational deliberation among the members of G in accordance with their individual and group epistemic normative requirements, would not result in further evidence that, when added to the bases of G 's members' beliefs that P , yields a total belief set that fails to make probable that P .⁴

The account has both summativist and non-summativist elements. Condition (i) makes group justification dependent on its members' justified beliefs, whereas condition (ii) adds the non-summativist idea that a justified group belief that P must survive a hypothetical deliberation of the group about its evidence on P . The first thing to notice is that Lackey's account, just like **Social Process Reliabilism**, focuses on the group members' doxastic, rather than propositional group justification. This can be fixed by replacing 'justifiedly believes' with 'has propositional justification to believe' **Group Epistemic Agent**. Second, the thusly corrected account still has an issue since it is too neutral. It requires that a significant percentage of the group members have justification to believe P . In cases where an epistemic advisor correctly recommends believing P to a less-informed advisee who isn't justified to believe P , only 50% of the group believes P , which would probably not be considered as 'significant'. Even if it was, in cases where we have not just one, but several advisees the percentage of group members with justification to believe P will be much lower than 50%.

Despite these problems with using **Group Epistemic Agent** for the purpose of precisifying **Collectivism**, it also provides some resources that can help with this. Condition (ii) states that group justification for believing P requires that if the members were to publicly share and discuss the evidence they all individually got, they wouldn't end up with a body of evidence that doesn't license believing P .⁵ Such a hypothetical public comparison of evidence could be a promising model of how to pool information-states in a truth-sided or veridical manner since it addresses the issues of how to handle inconsistent evidence. The proposal is that if a group has an (ideal) public discussion about its members' evidence, it

⁴This is, almost verbatim, Goldman's (2015: sec. 4.3) statement of Lackey's position.

⁵Hakli (2011) puts a similar constraint on group justification.

weeds out inconsistency in the evidence in favour of the better (true) evidence.

To sum up, recruiting the accounts of group justification currently on the market for the purpose of precisifying **Collectivism** proves difficult. One of the main reasons for this is that some of these accounts start with the group members' doxastic justification, which leaves out important evidence which actually determines what an advisee should collectively believe. However, Lackey's (ms) account might help with getting clearer on how to pool evidence. It is here where further research on **Collectivism** could benefit most from the group justification literature.

6.4 Is Epistemic Advice Factive?

6.4.1 The Factivity Necessity Condition

In this final section, I want to examine whether there is a factivity necessity condition (FNC) on correct epistemic advice. What I mean by this is whether 'You should believe P ' is true at a context of epistemic advice only if P is true. Note that this proposal is distinct from **Objectivism**, which I have rejected in the previous chapter. First, on **Objectivism**, 'You should believe P ' is true at a context of advice if (and only if) P is true. The proponent of an FNC, however, claims that truth of P is only necessary, not that it is also sufficient for 'You should believe P ' to be true. Second, the proposal that there is an FNC on correct epistemic advice does not have the absurd consequence of **Objectivism** that epistemic advice is only correct if the recommended doxastic attitude is belief or disbelief. The proposal does not apply to advice to have a doxastic attitude other than belief or disbelief since it only concerns sentences of the form ' S should believe P '.

I will argue that there is no FNC on correct epistemic advice. This insight is important for two reasons. First, it gives us more details on what a correct account of the truth-conditions of epistemic-advice sentences looks like (or better: does not look like). Second, as I explain in subsection 6.4.2, if one accepts an FNC, one should also accept a factivity *sufficiency* condition on correct epistemic advice. Since accepting such a condition and FNC would commit one to **Objectivism**, it is thus important for defending **Collectivism** to show that no FNC applies to

correct epistemic advice.

An FNC seems to be supported by intuition:

Police-False. Detective Lester Freamon is pondering about who killed Bob. His evidence is conflicting: it might either have been Omar or Marlo, but it must have been one of them. He looks at his colleague Kima Greggs, who has just entered the room. Unbeknownst to Lester, Kima has solid evidence that it was Omar: her reliable informant Bubbles told her that Marlo was at the other end of town at the time of the murder. However, Bubbles was mistaken, it wasn't Marlo who he saw at the other end of town. In fact, Marlo killed Bob.

(4) Kima: You should believe that it was Omar.

I believe many hear (4) as false because the belief that Kima recommends is false. Lester should not believe that it was Omar simply because he would then believe something false. Furthermore, if the aim of epistemic advice is to benefit the advisee epistemically, (4) must be bad advice since, if followed, the advisee would end up with something epistemically bad, i.e., a false belief. Assuming that bad advice is false advice, we can conclude that (4) is false.

I doubt that (4) is false, and I will present reasons for this in subsections 6.4.3 and 6.4.4. Before turning to this, let me quickly point out that **Collectivism** could account for an FNC. For example, we could specify it such that it becomes equivalent to **Group-Knowledge Collectivism**, which I mentioned in the previous section. This specification can be achieved by stipulating that the collective evidence is the counterfactual knowledge of the group, i.e., what the group would know if they believed it, and that believing a proposition P only accurately reflects an information-state i if P is fully supported by i .⁶ Since knowledge is factive, 'You should believe P ' is only true on this account if P is true.

⁶I have used the same method in subsection 3.4.1 to show that the knowledge norm can be integrated into **Doxastic Should**. There I also explain what I mean by 'full support'.

6.4.2 The Factivity Sufficiency Condition

As it turns out, our intuition does not only support that truth of P is necessary for ‘ S should believe P ’ to be true, but also that it is sufficient. Consider this case:

Police-True. Everything is as in **Police-False**. However, Kima is mischievous and wants to mislead Lester:

(5) Kima: You should believe that it was Marlo.

As we know, but Kima doesn’t, it was in fact Marlo. Just as (4) sounds false, (5) rings somewhat true. Lester should believe that it was Marlo; after all, it’s the truth. We can also give an argument for (5)’s truth that is analogous to the one we have considered for (5)’s falsity. If the aim of epistemic advice is to epistemically benefit the advisee, then uttering (5) is good epistemic advice since following it makes Lester adopt a true belief. Furthermore, since good epistemic advice is true, (5) is true.

Thus, there is not only intuitive support for an FNC, but also for a *factivity sufficiency condition* (FSC) on correct epistemic advice, i.e., the condition that ‘You should believe P ’ is true at a context of epistemic advice *if* P is true. But if there is both a factivity necessity and sufficiency condition on correct epistemic advice, i.e., if truth of P is both sufficient and necessary for ‘ S should believe P ’ to be true, we end up with the following claim:

Belief Objectivism. For any context of epistemic advice c_A , ‘ S should believe P ’ is true at c_A iff P is true.

Note that this is different from **Objectivism**:

Objectivism. For any context of epistemic advice c_A , ‘ S should have D to P ’ is true at c_A iff adopting D to P accurately reflects how much the fully realistic information-state supports P .

Objectivism entails that correct advice can only recommend belief or disbelief since only these are doxastic attitudes that are licensed by the fully realistic

information-state. **Belief Objectivism**, on the other hand, only applies to sentences of the form ‘*S* should believe *P*’ and does not give us the truth-conditions for epistemic-advice sentences recommending doxastic attitudes other than belief or disbelief.

Thus, someone can respect both intuitions supporting an FNC and an FSC on correct epistemic advice, without buying into seriously flawed **Objectivism**. However, **Belief Objectivism** is not sufficiently general because it does not provide the truth-conditions for epistemic advice concerning doxastic attitudes other than belief or disbelief. I think the best someone can do who accepts **Belief Objectivism** is to combine it with **Collectivism** in the following gerrymandered account:

Split. For any context of epistemic advice c_A , ‘*S* should have *D* to *P*’ is true at c_A iff (i) if $D =$ belief or disbelief that *P*, then having *D* to *P* accurately reflects how much the fully realistic information-state supports *P*, and (ii) if $D \neq$ belief or disbelief that *P*, then having *D* to *P* accurately reflects how much the collective evidence supports *P*.

As the reader can see, **Split** is different from **Collectivism** in that on **Collectivism** ‘*S* should believe *P*’ is not true iff *P*, but depending on whether the collective evidence licenses believing *P*. While **Collectivism** can be precisified such that the collective evidence can’t support a false proposition (see **Group-Knowledge Collectivism**), it can’t be precisified such that the collective evidence always supports the target proposition *P* if it is true. Unfortunately, we are often in situations where we lack evidence supporting what’s true.

Thus, if we want to respect the intuitions supporting an FNC and an FSC on correct epistemic advice, we are committed to **Belief Objectivism**. The best alternative to **Objectivism** if one wants one’s account to be consistent with **Belief Objectivism** is **Split**. This is a problem for **Collectivism** because it is incompatible with **Split**.

I see two main problems with **Split**. First, it breaks with any of the other accounts suggested so far on a very important point. All of these accounts present us with *one* advisory information-state, whereas **Split** states that there are two. On **Split**, there is thus technically not one ‘should’ of epistemic advice, but rather

two. One for epistemic advice to believe or disbelieve, relativized to the fully realistic information-state; one for epistemic advice to suspend judgement or to adopt a credence that doesn't correspond to belief or disbelief, relativized to the collective evidence.

What is unappealing about this is that it is at odds with the assumption that the pragmatic purpose of an utterance determines the contextually provided information-state. **Split** breaks with this since it entails that one kind of pragmatic purpose, epistemic advice, can determine two different kinds of information-states.

Now, the defender of **Split** could say that there is not one pragmatic purpose here, but two: that of epistemically advising to believe or disbelieve and that of epistemically advising to suspend judgement or adopt a credence. However, I think we are owed an explanation of why, first, in the case of epistemic advice the pragmatic purpose is broken up into sub-categories, and why, second, it is particularly broken up into belief/disbelief vs. suspending/credences. After all, the usual classification of doxastic attitudes is into coarse-grained (belief, disbelief, suspension) vs. fine-grained ones (credences).

Second, when constructing semantic theories, we often face a trade-off between simplicity and capturing our intuitions. Weatherson (2003) has argued that such a trade-off occurs between the JTB-analysis of 'knowledge' and an analysis that can account for Gettier cases. While the latter would validate more of our intuitions concerning knowledge-ascription, the JTB-analysis is simpler. He then draws on David Lewis' theory of meaning, according to which the meanings of terms must be both natural, which includes being simple, and must fit our use of the term, which is reflected in validating our intuitions. Weatherson argues that the JTB-analysis strikes a better balance between naturalness and use, i.e., simplicity and intuitive accuracy, than accounts sensitive to Gettier cases because the JTB-analysis is much simpler. Similarly, we can argue that **Split** is more intuitively accurate, but that **Collectivism** should, on balance, be preferred, since it is significantly simpler.

To sum up, **Split** is quite unattractive. This leaves us with the conundrum that our intuitions support **Belief Objectivism** but that there seems no plausible account that it is compatible with it. As a way out of this, I argue in the next subsection that these intuitions are mistaken and thereby defend **Collectivism**.

This argument should also satisfy those readers who might not have been very

convinced by the Weatherson-style argument and think that **Split** is better than **Collectivism**. If we have an account like **Collectivism** on the table, the only reason for going for a gerrymandered account like **Split** would be that the latter is more extensionally accurate since it predicts the truth-values of (4) and (5) correctly. This reason will be refuted if I can show that the intuitions supporting an FNC and FSC are mistaken.

6.4.3 Two Arguments against Factivity Intuitions

First, it would be ad hoc to postulate that ‘ S should believe P ’ is true iff P when it is used for epistemic advice, but that this isn’t the case when it is used for certain other purposes, even though all these purposes are pursued in circumstances of epistemic uncertainty. Take the following case of Skaiste, who is deliberating about what to believe with respect to proposition P :

(6) Skaiste: While I can’t be certain that P , the best explanation for phenomenon Y is P . Thus, I should believe P .

I want you to pay attention to the following things. First, (6) sounds absolutely fine. Second, it is the result of doxastic deliberation. Third, the ‘should’ in it is arguably subjective. Whether ‘I should believe P ’ is true in the context of (6) depends on whether Skaiste’s information-state supports P since we assume that there is no one she’s talking to, i.e., there is no collective evidence the ‘should’ in (6) could be relativized to. Finally, the sentence seems true despite the fact that Skaiste’s information-state does not entail P , but only abductively supports it.⁷ My point is that doxastic deliberation is done from a perspective of epistemic uncertainty. It often involves deciding what to believe about whether P in light of non-omniscient evidence, which allows both for P being true and false. In other words, responsibly forming a belief does not require infallibility. The upshot is that figuring out what one should believe about P in doxastic deliberation does not require figuring out what is supported by the fully realistic information-state.

Now, when giving epistemic advice, advisors are facing the same problems subjects who are doxastically deliberating face: epistemic uncertainty about whether the respective proposition P holds. True, there is (according to **Collectivism**)

⁷We’re assuming that what Skaiste says about her epistemic situation is correct.

a higher demand on them than on Skaiste: they need not just figure out what their personal evidence supports, but what the combined evidence of them and their advisee supports. However, such a joint body of evidence often leaves open the possibility that the proposition it supports is false. A defender of **Split** needs to explain why while both kind of agents, doxastic deliberators and epistemic advisors, are faced with epistemic uncertainty, only epistemic advisors need to know what's supported by the fully realistic information-state.

Second, as I have already pointed out in the discussion of **Objectivism** in the previous chapter, we seem to have intuitions about sentences of the form ‘*S* should believe *P*’ in contexts of epistemic advice even if we are not told whether the relevant target proposition is true or not. **Police** is an example for this. Just like **Objectivism**, **Split** has more difficulty explaining this than **Collectivism** does.

So far, I have not engaged with the arguments that (5) is true because it is good advice and (4) false because it is bad advice. In reply, I reject that (5) is good advice and (4) bad advice. In general, the idea that telling someone that he should believe *P* is good advice just as long as *P* is true and bad advice if *P* is false is too simplistic. Take the following analogy from football: a shot can be good, i.e., skilful, even if it does not achieve the ideal of football, that is, if the player doesn't score. Similarly, a shot can be bad or unskilful even if the player scores. I think that the same is true of advice. Even if I tell you to do *A* and if you did *A*, the aim of advice, that is, benefiting you, would be achieved, this can still be bad advice. In a case where our evidence does not support doing *A*, but doing *A* would actually benefit you, telling you to do *A* cannot be deemed good advice since only by luck does following the advice benefit you. Rather, giving you such advice is reckless of me. On the other hand, telling you to do *A*, when doing *A* won't benefit you, but our evidence clearly indicates that it will benefit you, is good, sensible advice, which unfortunately, without any fault of mine, will not benefit you.

6.4.4 Explaining Away Factivity Intuitions

While I might have shown that it's not the case that (4) is false in **Police-False** and (5) true in **Police-True**, I have not provided an *explanation* why our

intuitions are wrong in these cases.

Explaining away these intuitions is just an instance of solving a general puzzle that Doxastic Contextualism, but also contextualism about other terms such as other deontic modals or predicates of taste, have to solve. Take the following schematic case:⁸

Disagreement. Sam's information-state supports doing *A*. Sam utters:

(7) Sam: I should do *A*.

Unexpectedly, Nadia shows up. Her and Sam's evidence combined supports that Sam shouldn't do *A*. She utters:

(8) Nadia: You should not do *A*.

Nadia and Sam seem to be in disagreement. This is hard to explain on an information-sensitive contextualist semantics, at least given the traditional assumption that two speakers are in disagreement only if the sentences they uttered are inconsistent. According to contextualism, (7) and (8) are not inconsistent. The 'should' in (7) expresses the subjective sense and the 'should' in (8) the collective sense of 'should'.

Our intuition that (4) in **Police-False** is false and (5) in **Police-True** is true is just an instance of this general problem. To see this, consider that we cannot only get disagreement between two speakers in an (imagined) case, but also between a speaker in a case and a theorist evaluating the speaker's utterance. Consider again **Police-False**:

Police-False. Detective Lester Freamon is pondering about who killed Bob. His evidence is conflicting: it might either have been Omar or Marlo, but it must have been one of them. He looks at his colleague Kima Greggs, who has just entered the room. Unbeknownst to Lester, Kima has solid evidence that it was Omar: her reliable informant Bubbles told her that Marlo was at the other end of town at the time

⁸This is an abstraction of **Miners-unexpected**, a variation of the **Miners** case, which is described by Kolodny and MacFarlane (2010: 123). I will discuss it in more detail in subsection 7.3.1.

of the murder. However, Bubbles was mistaken, it wasn't Marlo who he saw at the other end of town. In fact, Marlo killed Bob.

(4) Kima: You should believe that it was Omar.

As mentioned, many would hear (4) as false. Imagine you belong to this group and would react somewhere along the following lines:

(T-4) You: (4) is false. Lester should not believe that it was Omar.

After all, it was Marlo.

Kima uttering (4) and you (T-4) is a form of disagreement.

In the next chapter, I will address the problem that disagreement causes for Doxastic Contextualism in more detail. I will argue, with Finlay (2014), that there is actually no disagreement where speakers are talking relative to different information-states and that they suffer from semantic blindness when they think so. When evaluating should-sentences, we sometimes make the mistake of evaluating the sentence not by evaluating the proposition actually expressed by it, but the proposition the sentence would express in our context. In cases such as **Disagreement**, a better-informed speaker like Nadia (sometimes) focuses on her own epistemic perspective, which makes her overlook the relativization of Sam's utterance to Sam's information-state. Focusing on what would be best for Sam, Nadia suffers from a "benign kind of confusion" (Finlay 2014: 241) and wrongly evaluates Sam's sentence 'I should do *A*' in light of her, better information-state.

I argue that the same goes on when someone hears (4) as false in **Police-False** and (5) as true in **Police-True**. Many, when presented with cases whose description include the truth-value of the target proposition, 'overlook' the relativization of doxastic-should sentences for the purpose of advice to the collective evidence, and go straight for an objective reading of (4) or (5). A reason for this could be that those with such intuitions have Lester's interest in true beliefs in mind, and thus focus on what would be a correct belief for him to hold, i.e., on what he should believe in the objective sense of 'should'.

Explaining our factivity intuitions away as a form of semantic confusion is therefore not ad hoc, but rather part and parcel of a comprehensive contextualist account of dealing with disagreement. Just as ordinary speakers get confused by the context-sensitivity of 'should', so do theorists like us.

One could object that while it might be plausible that ordinary speakers suffer from semantic confusion, theorists who are familiar with the context-sensitivity of ‘should’ should be educated enough to avoid this. I have two replies.

First, as Bealer (1998: 208) argues, one’s intuition that P can survive even if one comes to believe on basis of rational arguments that $\neg P$. Even though we know that the comprehension axiom of naive set theory is false, it still seems correct to us. Similarly, it can be argued that the trained semanticist or philosopher might still have the intuition that (4) is false in **Police-False**, even though she has been convinced that it is in fact true.

Second, our intuitions concerning the truth-values of should-sentences are actually quite sensitive to argumentation (Finlay 2014: 242). Take your statement (T-4). It seems easy to get you to retract this claim and agree that (4) is in fact true by confronting you with the following argument:

You’re right that it would have been best for Lester not to believe that it was Omar. But given Kima’s (and Lester’s) evidence, believing that it is Omar is certainly what they should do!

That the intuitions of some that (4) is false in **Police-False** and (5) true in **Police-True** are that easy to manipulate is further evidence that they are a result of semantic blindness.

6.5 Conclusion

In the previous chapter, I have argued for **Collectivism**, the view that epistemic advice of the form ‘You should adopt doxastic attitude D to proposition P ’ is true iff the combined evidence of the advisor and the advisee licenses D to P . In this chapter, I have examined **Collectivism** and its possible precisifications further. The first question I addressed is what we understand as the collective evidence. I have argued for a view of collective evidence that is either factive or at least truth-sided, i.e., that favours true evidence over false evidence (or more veridical over less veridical probability distributions) when the advisor’s and advisee’s evidence is pooled. Furthermore, I have expressed scepticism about utilizing accounts of group justification and knowledge for the purpose of making **Collectivism** more precise, with the exception of Lackey’s account of group

justification (ms). Finally, I have rejected the idea that epistemic advice of the form ‘You should believe P ’ is true only if P is true. This was important in order to defend **Collectivism** since respecting the intuition supporting such a factivity necessity condition on correct epistemic advice would push us to a view that is incompatible with **Collectivism**.

Chapter 7

The Relativist Challenge

7.1 Introduction

In this final substantial chapter, I present answers to relativist challenges to Doxastic Contextualism. These challenges are based on the kind of objections relativists direct at contextualist semantics for the information-sensitivity of the practical ‘should’ and of epistemic modals. As it is easier to understand their objections if one understands what relativism is, I begin this chapter with an explanation of the details of the view in section 7.2, focusing on MacFarlane’s (2014) very well-developed version of it.

Section 7.3 discusses a first difficulty for Doxastic Contextualism: the doxastic integration problem. It is analogous to the practical integration problem discussed by Kolodny and MacFarlane (2010), Björnsson and Finlay (2010), and Dowell (2013). The latter is a challenge to contextualist semantics for the information-sensitivity of the practical ‘should’. According to the practical integration problem, *Practical Contextualism*—as I call the view in short—struggles to explain how someone who is giving practical advice is addressing the question the advisee is deliberating about. The doxastic integration problem is the analogous problem for Doxastic Contextualism. I show that Dowell’s (2013) solution of the practical integration problem cannot be applied to Doxastic Contextualism. Instead, I defend Doxastic Contextualism by arguing that an advisor can be addressing the advisee’s question even if he is not answering a question with the same Kaplanian content, as long as it has the same Kaplanian character.

Section 7.4 moves on to the semantic assessment problem. It encompasses three different problems that all pertain to how speakers assess and interpret statements containing the doxastic ‘should’. First, the disagreement problem concerns how Doxastic Contextualism struggles to account for the apparent disagreement between speakers who both talk about what one of them should believe, but who, according to Doxastic Contextualism, actually say compatible things as they use different senses of ‘should’. I adapt Finlay’s (2014) solution to the analogous problem for Practical Contextualism, which suggests, on the one hand, that the disagreement actually does not lie in what is said, but rather in the speakers’ preferences, and, on the other hand, that on some occasions the speakers do not actually disagree, but merely suffer from semantic blindness (section 7.4.1). The second problem, the reporting problem, consists in the fact that we can accurately report statements about what someone should believe by disquoting them. Standard context-sensitive terms like indexicals do not share this feature of ‘disquotability’. However, as several authors have pointed out, there are other context-sensitive term, like ‘local’, that can also be disquoted. The disquotability of the doxastic ‘should’ therefore does not disqualify it from being context-sensitive (section 7.4.2). Finally, the retraction problem refers to cases where speakers retract their previous statements about what should be believed since they gathered new evidence in the meantime. According to Doxastic Contextualism, this would be inappropriate to do as the previous statements were true at the contexts at they were uttered. My defence is that this prediction of Doxastic Contextualism is, in fact, correct and that retracting those previous statements would be unnatural (section 7.4.3).

Of course, you don’t have to be a relativist to level these objections against Doxastic Contextualism, anybody can. The reason I am speaking of the ‘relativist challenge’ is that, first, these challenges were brought forward by relativists and, second, because relativists argue that the challenges speak in favour of their position since relativism does not face them. Another semantic account of the doxastic ‘should’ that does not face them is invariantism. This is the view that the doxastic ‘should’ is not context-sensitive, but always carries the same sense, for example, the objective or the subjective one, and that it is furthermore not information-sensitive in the way relativists say it is. Invariantism doesn’t have to deal with any of the mentioned problems since all of them stem from the fact

that Doxastic Contextualism assumes that the doxastic ‘should’ can change its senses in different contexts. Invariantism denies this. The reason I have chosen to discuss the challenges against the backdrop of relativism rather than invariantism is that chapters 3 to 5 were an all-round attack on invariantism. I have argued that the doxastic ‘should’ does not have one, but four different senses. If I have done my job, it should now be clear that invariantism is not a viable option. Only relativism can benefit from these challenges to Doxastic Contextualism.

7.2 Relativism

Relativist semantic theories have been defended for a variety of terms, such as epistemic modals like ‘might’ or ‘possible’¹, for ‘knows’², taste predicates³ and future contingents⁴. We can find rival contextualist proposals for epistemic modals⁵, ‘knows’⁶, and taste predicates⁷. As this list illustrates, the conflict between contextualism and relativism is not limited to the information-sensitivity of the deontic ‘should’ (or ‘ought’, for that matter), but extends to many other terms and has implications for the areas of philosophy that are concerned with the respective terms, such as aesthetics, epistemology, or metaphysics. I am solely concerned with the conflict between relativist and contextualist semantics for the information-sensitivity of the doxastic ‘should’ and do not intend to make any points about the other expressions.

The mentioned authors have somewhat different conceptions of what relativism is. Of the accounts of relativism on offer, I have chosen to focus on John MacFarlane’s. For one thing, his general theory of relativist semantics is extremely well-worked out (see his 2014 book). Furthermore, unlike all of the other relativists, he has provided a specific relativist semantics for the information-sensitivity of the deontic ‘should’.

As the reader knows, contextualism accounts for the information-sensitivity of the deontic ‘should’ by positing that ‘should’ has different semantic values

¹See Egan et al. (2005), Egan (2007), or MacFarlane (2011).

²See Richard (2004) or MacFarlane (2005).

³See Kölbel (2002), Lasersohn (2005), or MacFarlane (2014: Ch. 7).

⁴See MacFarlane (2003).

⁵See DeRose (1991), von Fintel and Gillies (2008), or Dowell (2011).

⁶See DeRose (1992) or Cohen (1998).

⁷See Glanzberg (2007).

depending on which information-state is provided by the context in which the sentence containing the term is uttered (the context of utterance). A sentence of the form ‘*S* should do *A*’ can, in one context, mean ‘*S* should do *A* in light of *S*’s information’, in another context mean ‘*S* should do *A* in light of the collective evidence’, and in yet another context mean ‘*S* should do *A* in light of the facts’. In other words, sentences of the form ‘*S* should do *A*’ express different propositions in contexts that provide different information-states.

On MacFarlane’s (2014: 286) truth-relativism, ‘*S* should do *A*’ expresses the same proposition at every context of use (provided, of course, that ‘*S*’ and ‘*A*’ have the same reference and ‘should’ has not changed its modal flavour). Information-sensitivity is rather modelled by the assumption that ‘*S* should do *A*’ has different truth-values relative to different context of *assessment*. A context of assessment is a possible situation in which someone can assess the truth of a sentence. This can be the same as the context of utterance—in cases where the speaker evaluates their own utterance—but often is not. An information-state is one of the parameters that a context of assessment can provide. A sentence of the form ‘*S* should do *A*’ is true at a context of assessment that provides the information-state *i* only if ‘*S* should do *A*’ is true in light of *i*.

At first sight, this seems to be a radical proposal. The truth-predicate is here not, as it is often assumed, a monadic predicate that just ascribes sentences (or propositions, beliefs, etc.) the properties truth or falsehood. It is rather taken to be a dyadic predicate that takes two arguments: a sentence *and* a context of assessment. Sentences are not true *simpliciter*, but only relative to a context of assessment. MacFarlane (2014: 49) claims that this is not as revisionary as it sounds. Standard semantic theories do not ascribe sentences truth-values *simpliciter*, but evaluate their truth relative to *circumstances of evaluation* (to use Kaplan’s (1989a, 1989b) terminology). For example, sentences are only true relative to possible worlds. The sentence

- (1) Barack Obama is the 44th President of the USA.

is true relative to the actual world, but not the world in which John McCain won the election in 2008. Some semantic models relativize sentences’ truth-values even further. On Kaplan’s (1989a, 1989b) model, sentences are evaluated at an index containing a world- and a time-parameter. A time-parameter is, on Kaplan’s view,

necessary in order to account for temporal operators like ‘yesterday’, which shift the time-parameter in the index. A sentence like ‘Yesterday I went to the shop’ is evaluated by evaluating the sentence under the scope of ‘yesterday’ relative to an index containing a world and whichever time ‘yesterday’ denotes.

To sum up, the truth-predicate of semantic theories is not monadic anyway. In standard semantic theorizing, sentences are at least evaluated relative to a possible-world-parameter, and in some cases even further parameters are added. What is new about MacFarlane’s truth-relativism are two aspects. First, it adds unusual parameters to the circumstance of evaluation, for example standards of taste or information-states. Second, these parameters are not restricted to parameters provided by the context of utterance.⁸ For example, we can think of a view on which the sentence ‘*S* should do *A*’ expresses, *ceteris paribus*, the same proposition at every context of utterance, but that the only information-state against which it can be evaluated is the one provided by the context of utterance. This view would agree with relativism that ‘*S* should do *A*’ is not true *simpliciter*, but would stipulate that it is only true relative to the information-state provided by the context of utterance.⁹ According to truth-relativism, by contrast, this information-state is provided by the context of assessment. Since the default information-state provided by the context of assessment is the assessor’s evidence, it is often *not* the information-state provided by the context of utterance. MacFarlane calls terms that are evaluated relative to unusual parameters provided by contexts of assessment ‘assessment-sensitive’.

To illustrate the difference between contextualism and relativism, let’s look at the **Miners** case, which I discussed first in subsection 2.5.1:

Miners. Ten miners are trapped either in shaft 1 or in shaft 2. Floodwaters threaten to flood the shafts. Sean has enough sandbags to block one shaft, but not both. If Sean blocks one shaft, all the water will go into the other shaft, killing any miners inside it. If Sean blocks neither shaft, both shafts will fill halfway with water, and just

⁸To be precise, this second condition is actually jointly sufficient and necessary for a semantics to be relativist in MacFarlane’s sense. The first condition is only *typical* of the semantics MacFarlane suggests for taste predicates, modals, etc. There is a possible relativist semantics that relativizes sentences only to a possible-world-parameter (see MacFarlane 2014: 89f.).

⁹MacFarlane (2009) calls this position ‘non-indexical contextualism’.

one miner, the lowest in the shaft where the miners are, will be killed.
Sean does not know in which shaft the miners are. He says:

(2) I should block neither shaft.

A physicist, who knows that the miners are in shaft 1, hears this and says to Sean:

(3) No, you should block shaft 1.

We will look at this case in more detail in the next section, but for now let's just focus on the following point. On a possible, and natural, contextualist reading of (2), Sean is speaking relative to his own information-state. That is, (2) is short for:

(2*) *In light of my [Sean's] evidence*, I should block neither shaft.

It follows that (2*), and thus (2), is true. Is (2) true according to relativism? In contrast to contextualism, there is no straightforward answer to this. Whether it's true depends on the context of assessment it is evaluated from. Relative to Sean's context of assessment, it is true because the information-state provided by his context of assessment is his own.¹⁰ This explains why Sean asserts (2). According to the pragmatics that MacFarlane (2014: Ch. 5) suggests for assessment-sensitive sentences, a speaker is permitted to utter such a sentence if it is true relative to their context (in this case context of utterance and assessment are the same). However, relative to the physicist's context of assessment (2) is false as the physicist's context of assessment provides her information-state. This explains why the physicist rejects Sean's utterance of (2).

After this introduction to MacFarlane's truth-relativism, I now turn to the first challenge to Doxastic Contextualism, the doxastic integration problem.

7.3 The Doxastic Integration Problem

The doxastic integration problem is analogous to the practical integration problem—as discussed by Kolodny and MacFarlane (2010), Björnsson and Finlay (2010), and

¹⁰However, as I'll explain in subsection 7.4.3 below, relativists are not committed to the view that the information-state provided by a context of assessment is *always* the assessors.

Dowell (2013)—which is that Practical Contextualism has difficulties explaining how practical advisors manage to address their advisees’ questions. In this section, I first introduce the practical integration problem and argue in agreement with Dowell (2013), and *pace* Björnsson and Finlay (2010), that it can be solved by *News-Sensitive Contextualism*, a version of Practical Contextualism (subsection 7.3.1). Subsequently, I show that analogous to the practical integration problem there is a doxastic integration problem, but that a doxastic version of News-Sensitive Contextualism provides no remedy for it (subsection 7.3.2). Finally, I present a solution to the doxastic integration problem (subsection 7.3.3).

7.3.1 The Practical Integration Problem

The practical integration problem, as it is dubbed by Björnsson and Finlay (2010), is a problem for Practical Contextualism. On this view, the semantics for the practical ‘should’ looks roughly as follows:

Practical Should For all contexts c that provide an ordering source g that reflects practical normativity, ‘Subject S should do A ’ is true at context c iff doing A -ing is the g -best available action for S in light of the information-state i provided by c .

The practical integration problem arises in cases where advice is given on what to do. Kolodny and MacFarlane (2010) introduce the **Miners** case to make this point. The case also brings out the semantic assessment problem, which I’ll deal with in the next section. Here I will only look at **Miners** with respect to the practical integration problem.

It seems clear that the physicist advises Sean on what to do. Now, Kolodny and MacFarlane (2010) claim that if the physicist advises Sean, ‘should’ needs to have the same sense in (2) and (3). Otherwise, the physicist would be “talking past” Sean (*ibid.*: 121). In this case, the physicist couldn’t be considered to be advising Sean since, as Björnsson and Finlay (2010: 11) put it, she would “not address the question over which [Sean] deliberates”. Similarly, Dowell (2013: 154) writes that (3) must be “an answer to the same question” that Sean is answering in (2).

One way to satisfy this constraint is for ‘should’ to be relativized to the same

information-state in (2) and (3). If the contexts in which (2) and (3) are uttered provide the same information-state, then (2) and (3) are answers to the same question.

Simple forms of contextualism struggle with this. For example, on a view we could call *Speaker Contextualism*, the information-state provided by the context is always that of the speaker. On such a view, Sean would be talking about what he should do in light of his evidence, while the physicist is talking about what Sean should do in light of her evidence. The physicist is talking past Sean.

A first stab at solving the practical integration problem is to suggest that the physicist is speaking to the collective evidence, that is, the pooled evidence of her and Sean. Since Sean is publicly deliberating about what to do, it is furthermore plausible that he's not intending to speak just relative to his own evidence, but to that of everyone he is speaking to, including the physicist's (Kolodny and MacFarlane 2010: 122). In this case, (2) is false. This can be made more palatable by arguing that it indeed seems hasty of Sean to conclude that he should block neither shaft if there are people around him who know where the miners are (Björnsson and Finlay 2010: 13).

Kolodny and MacFarlane (2010: 123) put pressure on such a view by giving a case in which the physicist arrives unexpectedly on the scene some time after Sean uttered (2) and then utters (3) (let's call this case **Miners-Unexpected**). Is the physicist part of the group whose evidence Sean was speaking relative to? Kolodny and MacFarlane (2010) argues that she is not. If she were, the evidence an agent would need to attend to in practical deliberation would become too far beyond the agent's reach to possibly play a role in deliberation. It would include even the evidence of people who are not present when one deliberates and are also not expected to turn up later. On the other hand, if the physicist is not part of the relevant group, the practical integration problem endures (Björnsson and Finlay 2010: 13).

Björnsson and Finlay (2010) and Dowell (2013) propose a view called *News-Sensitive Contextualism*, which is supposed to make the idea that the physicist's evidence belongs to the information-state Sean is speaking to more plausible. When we're deliberating about whether to do *A* or not, there is a certain time *t* by which we have to decide whether to do *A* or not. In **Miners**, *t* is the time by which the shafts are starting to flood and the first miner is risking to drown. Björnsson

and Finlay (2010) and Dowell (2013) suggest that what we're deliberating about is whether to do *A* in light of the information available to us by *t*. In this case, the physicist's evidence is part of the information-state Sean is speaking to and the practical integration problem is avoided. Furthermore, (2) is false, but reasonable, since Sean reasonably didn't expect to receive evidence by *t* about where the miners are.

Dowell (2013: 162) thinks News-Sensitive Contextualism solves the practical integration problem. Björnsson and Finlay (2010: 14, fn. 21), on the other hand, deny this. If News-Sensitive Contextualism is correct, then the physicist makes Sean's utterance of (2) false by making the evidence that the miners are in shaft 1 available to him. If she hadn't told him this, (2) would have been correct. The puzzle, for Björnsson and Finlay, is how the physicist is then helping Sean to "answer the question he is deliberating over" if Sean "would have been at least as likely to get the answer to the question right if she [the physicist] had chosen not to interfere" (ibid.:15).

Björnsson and Finlay suggest that this puzzle for Practical Contextualism is instead resolved by the following consideration: since deliberating agents have a preference for making decisions based on better information, Sean prefers an answer to the question of what he should do in light of evidence that includes the physicist's knowledge of the miners' whereabouts over an answer to the question of what he should do in light of evidence that doesn't include this information. The physicist is answering the more interesting first question, and can therefore be said to be advising Sean.

Björnsson and Finlay apparently accept that on this resolution of the puzzle, the physicist is not addressing the question Sean actually had in mind when uttering (2). They write: "agents' fundamental concerns in deliberating are [...] not simply to determine the correct answers to particular ought questions" (ibid.: 16). And further below:

"[I]f an advisor could make new and relevant information available to the agent, the proposition formerly of interest would lose [its] practical significance and become moot. Its significance passes to the question of what he ought to do relative to the new and improved body of information" (ibid.: 16).

Notice that they are therefore not just arguing that News-Sensitive Contextualism is not “sufficient to vindicate contextualism” (ibid.: 14, fn. 21), as they put it, but that they are effectively rejecting it. They seem to say that the physicist is addressing a different question than the one Sean had in mind, but that the physicist can nonetheless be said to be advising Sean.

This means that they are giving up on the constraint that advisors need to address their advisee’s original question. I think we can avoid this. A News-Sensitive Contextualist can argue that the physicist is addressing the very question Sean was deliberating about: ‘What should I do in light of the evidence available at the time t by which I have to act?’, where this evidence contains the physicist’s. It is just that if the physicist had not made her evidence available, this question would have meant ‘What should I do in light of E_{Hypo} ?’, where E_{Hypo} is the evidence that would have been available to Sean by t in this hypothetical situation, which would not have included the information that the miners are in shaft 1. However, the physicist actually makes the information available, and thus Sean’s question means ‘What should I do in light of E_{Actu} ?’, where E_{Actu} is the evidence that is actually available to Sean by t , which does include the information that the miners are in shaft 1. To conclude, the news-sensitive contextualist can maintain that the physicist *does* address Sean’s question. She helps him by making the question he is deliberating about the more informative question ‘What should Sean do in light of E_{Actu} ?’, rather than the question ‘What should Sean do in light of E_{Hypo} ?’, which it would have been, had she not uttered (3).

To conclude, I agree with Dowell (2013) that News-Sensitive Contextualism is capable of solving the practical integration problem. However, as I argue in the following, we can’t apply this solution to the doxastic integration problem.

7.3.2 Extending the Practical Integration Problem

The doxastic integration problem has not been discussed in the literature, but that it exists becomes clear if we look at a variant of **Police**, which was inspired by **Miners**:

Police-Disagreeing. Detective Lester Freamon is pondering about whether Omar or Marlo killed Bob, one of whom must have done it. His evidence is even-handed, it is equally likely on it that it was either

of them. He concludes:

(4) I should suspend judgement about who killed Bob.

His colleague Kima Greggs has solid evidence that it was Omar: her reliable informant Bubbles told her that Marlo was at the other end of town at the time of the murder. She hears Lester uttering (4) and says:

(5) No, you should believe that it was Omar.

On a contextualist semantics, the question Lester is answering by (4) is ‘Which doxastic attitude should I adopt to the proposition <It was Omar> in light of the information-state i_{c4} ?’, where i_{c4} is the information-state provided by the context in which he utters (4). (5), on the other hand, means ‘You should believe that it was Omar in light of i_{c5} ’, where i_{c5} is the information-state provided by the context in which Kima utters (5). The *doxastic* integration problem is that Kima seems to be giving advice to Lester, but that it is hard for contextualism to explain how Kima is addressing the question Lester is deliberating about.

One might think that a doxastic equivalent to News-Sensitive Contextualism is motivated by the following variant of **Police**:

Police-Unexpected. Everything is as it is in **Police-Disagreeing**, except that Kima is not present when Lester utters (4) at t_4 . He also does not expect her to later contribute anything to the case. At t_5 , Kima shows up unexpectedly, having collected the same evidence as in **Police**. She utters (5).

The idea would be that if we want Kima to address the question Lester was thinking about, we need him to speak about the information-state that will be available to him at or after t_5 . *Doxastic News-Sensitive Contextualism* would try to make this idea plausible by suggesting that when we are deliberating doxastically, we are deliberating about which doxastic attitude we should adopt to the relevant proposition P in light of the evidence that will be available to us by the time we need to form a doxastic attitude to P . I reject Doxastic News-Sensitive Contextualism for three reasons.

First, such a view needs to determine what the time is by which one has to form a doxastic attitude. For certain propositions, there is no such time. These include propositions that are of purely theoretical interest, which have no relevance for practical purposes, and irrelevant propositions, e.g., the proposition that the number of stars in the universe is even.

Second, there is a crucial asymmetry between the doxastic and the practical ‘should’, which makes Doxastic News-Sensitive Contextualism implausible. When we are deliberating about what we practically should do, we’re often concerned with the question of which of a set of possible actions we should perform in the future (rather than right now). These are the kind of cases that motivate *Practical* News-Sensitive Contextualism: where the time of deliberation and the time of action are not identical. However, when we are deliberating about what we doxastically should do, we never really seem to be concerned with which doxastic attitude we should adopt in the future. Rather, we try to figure out which doxastic attitude to adopt right now, at the time of deliberation. That is, it looks like the time of deliberation and the time of ‘action’, i.e., of forming the relevant doxastic attitude, are identical.

This asymmetry is crucial since it might be plausible to deliberate about what one should do at a future time in light of the information one will have at that time; it is absurd, however, to deliberate about which doxastic attitude we should form *now* in light of the information we will have at some point in the future. One reason why this is absurd is that it is never too late to change your mind. The doxastic attitude to *P* you form now, you can easily revise at the time *t* when you have to act on it, in light of the evidence you have then. Why would you then deliberate about which doxastic attitude you should adopt *now* in light of this future information-state?

Third, Doxastic News-Sensitive Contextualism fails a test that its practical counterpart passes. If Practical News-Sensitive Contextualism is correct, then Sean should react as follows if he is told that he will receive evidence about where the miners are before the time he has to act:

- (6) OK, I should not block either shaft then, but I don’t know yet which of the two I should block. I will need to wait for the information about where the miners are.

This sounds like the correct response. It supports Practical News-Sensitive Contextualism since Sean obviously speaks about what he should do in light of the evidence that will be accessible to him in the future, not the one that is accessible to him now.

On the other hand, if we tell Lester that in an hour, Kima will tell him who killed Bob, it would sound infelicitous if he replied:

- (7) # OK, I shouldn't suspend judgement now then, because in light of Kima's evidence I should either believe or disbelieve that it was Omar.

The appropriate response seems rather to be:

- (8) I should suspend judgement on whether it was Omar for now then.

If Doxastic News-Sensitive Contextualism was correct, (7) would be true and reasonable, but (8) false and unreasonable. However, it seems to be exactly the other way around.

I conclude that Doxastic News-Sensitive Contextualism is not a solution to the doxastic integration problem.

7.3.3 Solving the Doxastic Integration Problem

I suggest that we address the doxastic integration problem as follows. First, with respect to **Police-Disagreeing** I propose a rather straightforward solution. In chapter 5, I have put forward the following view:

Collectivism. For any context of epistemic advice c_A , ' S should have D to P ' is true at c_A iff adopting D to P accurately reflects how much the collective evidence supports P .

On **Collectivism**, Kima is speaking relative to her and Lester's collective evidence. This helps with respect to **Police-Disagreeing** if we assume that Lester is speaking relative to the same body of evidence. This is plausible. As Björnsson and Finlay (2010: 13) put it with respect to the practical 'should', where someone

is deliberating publicly, that is, in a dialogue with others, they arguably see their deliberation “as a shared problem, to be solved with collective resources.” On this account, (4) is false, but we can give a good explanation for it. It is hasty for Lester to conclude that he should suspend judgement on whether it was Omar, if he could simply ask Kima what her evidence concerning the case is.

The story is a bit more complicated concerning **Police-Unexpected**. Since Lester is not publicly deliberating—but rather privately on his own—we need to assume that by uttering (4) he is answering the question of what he should believe, in light of his own evidence, about who killed Bob.¹¹ When Kima appears and utters (5), she talks about what he should believe in light of their collective evidence. It therefore seems like she is addressing another question.¹²

My main claim is that in *some sense* she is addressing the same question: They are both trying to answer the question ‘What should Lester believe about who killed Bob?’ True, according to Doxastic Contextualism this answer has different contents if asked in Lester’s and Kima’s contexts respectively. Nonetheless, they are in some sense the same.

Let me explain. Consider the following analogy: When Ann and Boris are respectively asking ‘Who am I?’, they are in some sense asking the same question. In some other sense, they are not. Ann’s question means ‘Who is Ann?’ and Boris’ question means ‘Who is Boris?’. They are the same questions in the sense that their Kaplanian (1989a, 1989b) character is identical; they are different questions in the sense that their Kaplanian content is different. The character of ‘Who am I?’ is a function from contexts to contents, where the assigned content changes with the speaker of the respective context. Let’s distinguish accordingly between the character and the content of a question.

The question Kima is addressing has the same character as Lester’s.¹³ The character of ‘What should Lester believe about who killed Bob?’ is a function

¹¹More precisely, he is answering the question of which doxastic attitudes he should adopt towards the propositions <Omar killed Bob> and <Marlo killed Bob>. I use ‘what he should believe’ as a short form for this.

¹²For more on the distinction between private and public doxastic deliberation, see section 5.2.

¹³This is not entirely correct. Since Lester formulates the question in the first person, his question is ‘What should *I* ...?’, whereas Kima’s question is ‘What ... should *Lester*...?’. Because Lester’s question contains the indexical ‘I’ and Kima’s the proper name ‘Lester’, the questions actually do not have the same character. However, this difference in character is irrelevant to the point I’m making.

from contexts to contents, where the content differs with the information-state that is provided by the respective context.

When Ann tells Boris the answer to her question ‘Who am I?’, she is hardly helping him with his question ‘Who am I?’ Thus, one might think that the fact that Kima is addressing a question which has the same character as Lester’s doesn’t help solving the doxastic integration problem. However, there are cases where addressing a question with the same character, but different content, seems to help:

Score. Carlo is watching a football match between Borussia Dortmund and Bayern München. At 8 pm, the score is 0:0. His friend Diana calls him at 8 pm and asks:

(9) ‘What is the score?’.

Before Carlo can reply, the telephone connection is interrupted. At 8.05 pm, Aubameyang scores for Dortmund. Carlo calls Diana and says:

(10) ‘Hey! To answer your question: The score is 1:0.’

(10) sounds perfectly fine; Carlo seems to answer Diana’s question expressed by (9). However, notice that there is a hidden indexicality in (9). (9) means ‘What’s the score *now*?’ Carlo is answering the content this question has at the context of his utterance of (10), that is, he’s answering ‘What is the score at 8.05 pm?’ However, the content of (9) at the context where Diana uttered it is ‘What is the score at 8 pm?’ Nonetheless, (10) sounds right: Carlo is answering, at least in some sense, Diana’s question.

To conclude, at least in some cases it looks like speakers are answering the question of their addressee when they are answering the content of a question that has the same character as the addressee’s question, but a different content. I suggest that questions like ‘What should I believe about who killed Bob?’ are in this respect more similar to questions like ‘What’s the score?’ than they are similar to questions like ‘Who am I?’

One reason to think this is that it simply *seems* like Kima is answering Lester’s question. And as **Score** shows, the mere fact that there is an indexical element

in the question is not necessarily a theoretical reason undermining this intuition. Furthermore, **Score** and **Police-Unexpected** share a property that makes it likely that they are alike in the mentioned respect. Diana is more interested in an answer to the content of her question relative to the context of 8.05 pm rather than relative to the context of 8 pm. That is, she wants to know what the score is at 8.05 pm, rather than what the score is at 8 pm. Similarly, Lester is more interested in an answer to the content of his question at the context where Kima utters (5) than at the context where Lester utters (4). Thus, where a speaker's and a subject's questions have the same character but not the same content, the speaker might be taken to answer the subject's question if the subject is more interested in an answer to the content of the speaker's question rather than their own. This also explains why it sounds wrong to say that Ann addresses Boris' question 'Who am I?' by telling him the answer to the question 'Who is Ann?' Boris is not interested in an answer to this question, or at least not more than in an answer to the question 'Who is Boris?'

Notice that this is the same explanation as the one Björnsson and Finlay (2010) offer for how it is possible that the physicist can, legitimately, address a question that is not the one Sean originally asked, but nonetheless be said to advise Sean. The difference between us is only that Björnsson and Finlay (2010) concede that the physicist is answering a different question, full stop. I, on the other hand, argue that there is no need to make such a concession. The physicist is answering the same question in the sense that it has the same character. In the end, Björnsson and Finlay (2010) would probably be happy to accept this improvement on their account.

Contra Björnsson and Finlay (2010), I argued that Practical News-Sensitive Contextualism can solve the practical integration problem. However, the same reasoning I just outlined with respect to the doxastic integration problem can arguably also be applied to solve the practical integration problem. I will leave the question of whether the latter is to be preferred to Practical News-Sensitive Contextualism to future research and turn now to the second challenge to Doxastic Contextualism.

7.4 The Semantic Assessment Problem

The semantic assessment problem actually consists in a set of problems that concern how speakers semantically evaluate and interpret their own and each other's statements about what should be done and believed. More specifically, contextualist accounts of the deontic 'should' have difficulties explaining the following three phenomena: First, it seems that in cases like **Miners** and **Police** advisors and advisees can *disagree*. Second, it looks as if speakers can successfully *report* and refer to statements about what should be done and believed. Third, it appears appropriate for speakers who later on receive more relevant information to *retract* their previous statements of this sort.

I will attend to these problems in turn in the following. All of them apply equally to Practical and Doxastic Contextualism and hence almost all of the solutions that I present solve the respective problem for both views. Furthermore, the problems are not restricted to contextualist semantics for the deontic 'should', but have received much attention in the literature on the semantics for epistemic modals. I will make reference to this literature where it helps clarity.

7.4.1 Disagreement

Remember, in **Miners**, Sean comes to conclude:

- (2) I should block neither shaft.

This is countered by a physicist, who knows that the miners are in shaft 1 and who says to Sean:

- (3) No, you should block shaft 1.

The practical integration problem arises since the physicist is advising Sean. The problem we are concerned with here, let's call it the *disagreement problem*, arises because it seems as if the physicist and Sean are disagreeing.¹⁴ On a natural picture, disagreement occurs where one agent accepts a proposition P and another agent rejects P . This presents a simple form of contextualism, on which the contextually provided information-state is always the speaker's, with a difficulty. According to it, (2) means

¹⁴I already touched on the issue of disagreement in subsection 6.4.4.

(2*) In light of my [Sean's] evidence, I should block neither shaft.

and (3) means

(3*) In light of my [the physicist's] evidence, you should block shaft 1.

Thus, that Sean accepts (2) doesn't mean that he rejects (3) and that the physicist accepts (3) doesn't mean that he rejects (2). Doxastic Contextualism has the same problem as Lester and Kima appear to be disagreeing with each other.

Kolodny and MacFarlane (2010: 143) point to this disagreement data as evidence in favour of relativism. According to relativism, two parties disagree if (i) they respectively accept and reject a proposition P and (ii) accepting and rejecting P aren't simultaneously accurate relative to any context of assessment (MacFarlane 2007: 26f.). Since Sean accepts and the physicist rejects the proposition <Sean should block neither shaft>, and either accepting *or* rejecting this is accurate at a context of assessment (depending on which information-state is provided by this context), but not both, Sean and the physicist are disagreeing.

One way to solve this problem for contextualism is to stipulate that Sean and the physicist (or Lester and Kima) are in fact speaking relative to the same information-state. This is the strategy Dowell (2013) chooses by advocating Practical News-Sensitive Contextualism. Both in **Miners** and **Miners-Unexpected**, Sean and the physicist are talking about the same issue: whether Sean should block a shaft in light of the evidence that will be available to him by the time he has to decide. Sean accepts this, while the physicist rejects it, so they are disagreeing. Unfortunately, News-Sensitive Contextualism is not an adequate model for the doxastic 'should', so at least with respect to the doxastic 'should' it won't help solving the disagreement problem.

An alternative way to address the disagreement problem is to argue that there can be disagreement between the parties, even though there is no proposition that one agent accepts and the other rejects. Maybe the disagreement is not due to a conflict between beliefs, but rather due to a conflict between preferences (Finlay 2014: 219). By uttering (2), Sean expresses a preference for blocking neither shaft, whereas the physicist express a preference for him to block shaft 1. Similarly, Lester is expressing a preference for withholding judgement on who killed Bob, whereas Kima expresses a preference for Lester to believe that Omar killed Bob.

Sean and the physicist, as well as Kima and Lester, are accordingly disagreeing since they have conflicting preferences.¹⁵

The limits of this approach lie with scenarios in which speakers explicitly refer to and deny what another agent has said. For example, the physicist could just as well say in response to Sean:

(11) That is false. You should block shaft 1.

Or:

(12) What you said is false. You should block shaft 1.

In the first sentence of (11) and even more in the first sentence of (12), the physicist is explicitly referring to and denying what Sean said. When the physicist responds to Sean in this manner, it is hard to defend that the disagreement between Sean and the physicist is purely one of preferences. It is not just that they *want* things that are inconsistent and can't both be obtained, but the physicist's utterance implies that they also *say* (and hence believe) inconsistent things (Finlay 2014: 239).

In their original defence of Practical Contextualism, Björnsson and Finlay (2010) propose a pragmatic solution to the disagreement problem, which acknowledges that in **Miners-Unexpected** (2) and (3) actually do not express inconsistent propositions since Sean and the physicist are speaking relative to different information-states. They claim that the first sentences of (11) and (12) are, nonetheless, not false. The physicist is not referring to the proposition Sean actually expressed when he uttered (2), but rather to the proposition he would express if he uttered (2) in the physicist's context. If (2) was uttered in the physicist's context, it would express that in light of the collective evidence Sean should block neither shaft, which is wrong and is what the physicist is denying. Björnsson and Finlay (2010: 22) argue that the physicist is correctly referring to this proposition because should-sentences pragmatically function as recommendations. Since recommendations should be checked against the best available

¹⁵Stevenson (1944) proposed this idea of a 'disagreement in attitude' as part of his emotivist account of moral language. Huvenes has mobilized the idea to defend contextualism about taste predicates (2012) and epistemic modals (2015) in the face of disagreement data.

information, the physicist accurately evaluates (2) as uttered in her context rather than Sean's context, as it provides a better-informed information-state.

Finlay (2014: 240f.) has come to reject this proposal as “too cavalier about the role of the demonstrative ‘that’, or talk about what the speaker ‘said’ or ‘believed’”. I think he's right. It seems absurd to claim that by ‘what you said’ Sean could refer to a proposition that Sean did not assert, i.e., precisely something that Sean did *not* say. Even with respect to ‘that’ it appears clear that it must pick out what Sean actually said.

Instead of this pragmatic solution, Finlay (2014) prefers to attribute semantic blindness to speakers. The physicist in fact refers to the proposition that Sean expressed and thus errs in her assessment of Sean's utterance as false. This does not mean to give up on the idea that there is disagreement between Sean and the physicist (or Lester and Kima), since it still holds that their preferences conflict. However, it admits that there is no conflict in their beliefs and that the physicist mistakenly thinks so. I adopt Finlay's (2014) suggestion here. If Kima were to say

(13) What you said is false, Lester. You should believe that it was Omar.

the first sentence of this would be false.

What are the theoretical costs of positing semantic blindness? Finlay (2014: 241) says they are not too high and calls mistakes like Sean's a “benign kind of confusion.” This confusion arises for two reasons. First, if you assess a should-sentence that was uttered in a context that provides a different information-state than yours and then want to give advice on what to do (or to believe), you have to attend to two different perspectives at (almost) the same time, which is cognitively challenging. This is certainly right. Advisors are, naturally, focused on giving advice and therefore attend to the collective evidence. This fixation makes a shift to their advisee's evidence in order to assess their advisee's utterance harder. Second, ordinary speakers might have difficulties tracking the difference between disagreements due to a conflict between beliefs and those due to a conflict between preferences. Disagreement in belief is expressed via evaluations like ‘That/what you said is false’. The physicist might have detected correctly that there is disagreement between her and Sean, and then chosen a phrase we often

use to express disagreement, overlooking that it's not appropriate for the kind of disagreement at hand.

Not only can we explain the semantic blindness of speakers like the physicist, it also looks like that such speakers are quite willing to admit to a mistake if pressed (ibid.: 243). Again, imagine that the physicist explicitly disagrees with Sean in **Miners-Unexpected**:

(12) What you said is false. You you should block shaft 1.

Now, consider someone overhears the physicist and approaches her:

(14) Actually, Sean was speaking truly. He must have been talking about what he should do in light of *his* evidence, and this is certainly to block neither shaft.

It is very plausible that the physicist would react as follows:

(15) Fair enough, given his evidence he came to the right conclusion. But given that I know that the miners are in shaft 1, he should block shaft 1!

The physicist's retraction of her earlier statement sounds perfectly natural. Therefore, attributing semantic blindness to her when she utters (12) is probably exactly what we should do. We can easily imagine Kima retracting (13) if someone challenged her in a similar manner by pointing out that Lester must have been talking to his own evidence. To sum up, the disagreement data can be explained in accordance with Practical and Doxastic Contextualism by reference to the phenomenon of disagreement between preferences and the phenomenon of semantic blindness.

7.4.2 Reports

We have already touched upon the issue of said-that-reports in the previous subsection. I agreed with Finlay (2014) that speakers at times get confused by the information-sensitivity of the deontic 'should' and go wrong when trying to evaluate what others said about what should be done. In this subsection, I will address the question of how, on a contextualist account of the information-sensitivity of the deontic 'should', it is possible at all to report what others say

or believe about what should be done or believed. Let's dub this the *reporting problem*.

Miners-Unexpected helpfully illustrate the problem. Here is another possibility of how the physicist could choose to assess Sean's utterance:

(16) You said that you should block neither shaft. But this is false.
You should block shaft 1.

The main difference between (12) and (16) is that in (16) the physicist reports what Sean said by saying 'You said that you should block neither shaft' instead of just referring to it with 'what you said.' As it turns out, it is actually tricky for contextualism to explain how the first sentence of (16) can accurately report what Sean said. The information-state provided by the physicist's context is the collective evidence. So it suggests itself that this is the information-state that 'should' in the first sentence of (16) is relativized to. The sentence would then mean:

(16*) You said that, in light of the collective evidence, you should block neither shaft.

However, this fails to report what Sean said. Rather, what we want it to mean is:

(16**) You said that, in light of your evidence, you should block neither shaft.

Egan et al. (2005: 138) point to this phenomenon with respect to epistemic modals. When epistemic modals are embedded in a belief-that- or said-that-construction, the information-state they are relativized to shifts from the one originally provided by the context of the report to the information-state provided by the context of the reported speech act or belief.

Finlay (2014: 237) claims that this makes sense since "reporting a person's speech or thought makes her content or perspective salient." While a contextualist semantics for the information-sensitivity of the deontic 'should' indeed requires this, it is far from clear that this is unproblematic. What Finlay does not mention is that such behaviour cannot be observed in standard context-sensitive terms like indexicals (Egan et al. 2005: 148):

(17) Tim: I'm hungry

(18) Görkem: Tim said I'm hungry.

When 'I' or other indexicals like 'now' or 'here' are embedded in a said-that-construction, the contextual parameter fixing their reference does not shift away from the default one. The reference of 'I' in (18) is the speaker of the context of utterance, Görkem, and does *not* shift to Tim, even though Tim is the speaker of the context of the speech act that Görkem is reporting. In other words, indexicals cannot be disquoted in said-that-reports across contexts, whereas modals like 'might' or 'should' can.

Egan et al. (2005: 149) claim that this speaks against the context-sensitivity of epistemic modals. If uncontroversially context-sensitive terms cannot be disquoted, but epistemic modals can, why should we assume that epistemic modals are context-sensitive? Obviously, the same argument works for deontic modals.¹⁶ Notice that a relativist semantics can avoid the reporting problem altogether. As sentences of the form '*S* should do *A*' (or 'It might be that *P*') do not, *ceteris paribus*, change their content across contexts on relativism, there is no question how, on this semantics, the first sentence of (16) could accurately report what Sean said.

However, there are other context-sensitive term, such as 'local', 'nearby', or 'enemy', which share this property with epistemic and deontic modals. Disquotability does therefore actually not undermine contextualism.¹⁷ While Egan et al. (2005) even cite 'local' as *support* for their position, we'll see that it in fact works against them.

'Local' is a plausible candidate for context-sensitivity since if I speak of a 'local bar' while I am in Edinburgh and you speak of a 'local bar' while you're in St Andrews, I mean 'a bar in Edinburgh' and you mean 'a bar in St Andrews'. That is, the content of 'local' varies with the location of the context. Egan et al. (2005: 148) present the following dialogue as evidence that 'local' cannot be disquoted:

¹⁶To be clear, Egan et al.'s (2005) argument is only directed at the claim that the information-sensitivity of epistemic modals is a form of context-sensitivity. They do not challenge that the different flavours of modality of modals, i.e., epistemic, deontic, circumstantial, are a form of context-sensitivity.

¹⁷Cappelen and Lepore (2006: 475), Humberstone (2006: 315), and Cappelen and Hawthorne (2009: 39f.) have made this point.

Local Bar. Brian is calling from Providence, Hud and Andy are in Bellingham.

(19) Brian: When I get all this work done, I'll head off to a local bar for some drinks.

(20) Andy: How much work is there?

(21) Brian: Not much. I should get to the bar in a couple of hours.

(22) Hud: Hey, is Brian in town? Where's he going tonight?

(23) # Andy: He thinks he'll be at a local bar in a couple of hours.

I agree that Andy fails to report Brian's utterance of (21) in his utterance of (23). There is only one plausible way in which Hud could interpret Andy: that Brian said that he'll be in a local bar in Bellingham. However, notice that the case is set up in a very peculiar way to achieve this effect. Hud is asking Andy whether Brian will be "in town", i.e., in Bellingham¹⁸. Andy does not deny this and then says that Brian said he'll be in a local bar in a few hours. By not denying, and therefore implicitly *confirming* that Brian will be in Bellingham later, Andy seems to intend to communicate with (23) that Brian will be in a bar in Bellingham later.

Things look quite different if a case like this is set up differently:¹⁹

Concerned Parents. Arun is calling his parents from Manchester, where he is going to university. His parents, Dhara and Raj, live in London.

(24) Dhara: How are things? Do you have any plans tonight?

(25) Arun: I'll be going to a local bar later with some of my friends.

(26) Raj: So, how is our son in Manchester?

(27) Dhara: Oh, he seems fine. He said he'll be going to a local bar later tonight.

¹⁸For the non-American reader: Bellingham is the largest city, and the county seat, of Whatcom County in the State of Washington, USA.

¹⁹Humberstone (2006: 315f.) describes a very similar case.

The last sentence of (27) sounds absolutely fine and seems to perfectly report Arun's plans. It appears that the location parameter is switched from Dhara's location to Arun's and 'local bar' therefore comes to mean 'bar in Manchester'. What did the trick was that Raj made 'Manchester' the salient location in the conversation, which Dhara's use of 'local' could feed on.

To conclude, a presumably context-sensitive term like 'local' can, at least on some occasions, be disquoted. Accordingly, the fact that modals can be disquoted should not be counted as clear evidence against their context-sensitivity as Egan et al. (2005) argue. Admittedly, I struggle to think of any case where a deontic 'should' cannot be disquoted due to its information-sensitivity, whereas 'local' can at least not be disquoted in Andy's utterance of (23). However, once we see that some context-sensitive terms are only *sometimes* incapable of being disquoted across contexts, it does not seem that implausible that some context-sensitive terms are *never* incapable of disquotation across contexts. If there is other evidence supporting the term's context-sensitivity, its 'disquotability' should not bother us too much. That speakers actually seem to be aware of the context-sensitivity of the deontic 'should', as for example the physicist is when she retracts her previous judgement of Sean's utterance in (15), is such evidence.

7.4.3 Retraction

Retraction occurs when a speaker withdraws an earlier statement and declares that they have been wrong. Björnsson and Finlay (2010: 18, n. 26) describe the following case as the kind of data relativists could point to:

Miners-Retract. Everything is as in **Miners** (p. 177), except that after Sean comes to conclude that he should block neither shaft and states (2), the physicist does not show up. Rather, shortly before the shafts get flooded, Sean finds out himself that the miners are in shaft 1. He says:

(28) I was wrong. I should block shaft 1.

The idea is that (28) sounds right, but that contextualism can't account for this if it posits that Sean is talking relative to different information-states when he's stating (2) and (28). If this were the case, Sean should not retract what he said

earlier since blocking neither shaft is what he should have done in light of the evidence he had at the time. Again, relativism has no difficulties with this. When Sean utters (2), (2) is true relative to the information-state of this context of utterance. When he later retracts it, he does so in a new context of assessment, whose information-state has been updated with the whereabouts of the miners, such that (2) is false relative to this new context, which makes the retraction appropriate.

I only know of Björnsson and Finlay (2010) who consider retraction data as a potential threat to their Practical Contextualism, and am not aware of any relativist who actually uses retraction data to argue for relativism about deontic modals. However, MacFarlane (2011: 148) uses data of this sort to motivate relativism about *epistemic* modals, and it is thus worthwhile to discuss this *retraction problem* for Doxastic and Practical Contextualism. As it turns out though, it seems to me the least worrisome of our three problems. First, it is important to notice that (28) is uttered before the shafts get flooded. While (28) already sounds a bit off to me, (29) in the following variant of the case definitely does:

Miners-Post. Everything is as in **Miners-Retract**, except that Sean only finds out that the miners were in shaft 1 *after* the shafts have been flooded and one of the miners died.

(29) I was wrong. I should have blocked shaft 1.

Björnsson and Finlay (2010: 18, n. 26) write that such post-mortem retractions are not “as natural” as the retraction in **Miners-Retract**. I would say that they aren’t natural at all. Clearly, Sean is too hard on himself. His original judgement was correct, given the information he had. In fact, blocking shaft 1 would have been irresponsible given that he would have risked letting ten people die.²⁰

This leaves us with **Miners-Retract**. As the reader might have noticed, the case does not cause difficulties for Practical News-Sensitive Contextualism. (28) is uttered prior to the time when Sean has to decide whether to block one of the shafts or not. On Practical News-Sensitive Contextualism, (2) and (28) are therefore

²⁰In line with this, von Fintel and Gillies (2008) argue that retractions of statements containing epistemic modals are sometimes not very natural either, even though predicted by relativism.

spoken relative to the same information-state, namely the evidence available to Sean by the time he needs to act. And relative to this information-state, (2) was in fact wrong.

This is not to endorse Practical News-Sensitive Contextualism, but it shows that the contextualist who wants to treat (28) as true has options. However, as I have argued in section 7.3.2, News-Sensitive Contextualism does not work for the doxastic ‘should’. Thus, I am committed to the view that whenever someone correctly changes their previous judgement of what they should believe based on new evidence they have gathered, they cannot reasonably retract this previous judgement since it was true relative to the impoverished evidence they had at the time. I do not see this as a bug, but rather as a feature of my view. Consider the following variant of **Police**:

Police-Retract. Lester is pondering about whether Omar or Marlo killed Bob, one of whom must have done it. His evidence is even-handed, it is equally likely on it that it was Omar as well that it wasn’t. He concludes:

(4) I should suspend judgement about who killed Bob.

Later that day, Kima tells him that Bubbles testified that he saw Marlo at the other end of town at the time of the murder. Lester reacts to this as follows:

(30) I was wrong. I should have believed that it was Omar.

Lester’s reaction is mistaken. He was not wrong when he uttered (4), suspending judgement is what he was supposed to do when his evidence was even-handed.

One might think that **Police-Retract** is even evidence *against* relativism since the relativist needs to claim that (30) is correct given that relative to the information-state of Lester’s context of assessment (4) is false. This is not the case though. MacFarlane (2014: 298) remarks that his model allows that the information-state provided by a context of assessment does not always have to be the evidence of the assessor at the time of assessing. When one is looking at one’s own or others’ past doxastic attitudes, for example, the information-state

can shift to the evidence of the respective person at the time. **Police-Retract** hence favours neither contextualism nor relativism.

7.5 Conclusion

In this chapter, I have defended Doxastic Contextualism against two potential relativist challenges. The first challenge was the doxastic integration problem, which I solved by pointing out that an epistemic advisor can be seen as addressing the question of their deliberating advisee even if the advisee and the advisor are speaking relative to different information-states. The trick is that even though they are concerned with questions that have a different Kaplanian content, their questions share the same Kaplanian character.

The second challenge was the semantic assessment problem, which consists of three different problems. The retraction ‘problem’ is not much of a problem at all as I dispute that it is felicitous for speakers to retract previous judgements after their evidence has changed. The reporting problem is more serious since it points to a semantic difference between standard context-sensitive terms like indexicals and the doxastic ‘should’, namely the latter’s disquotability. As we have seen, however, other context-sensitive terms, like ‘local’, can also sometimes be disquoted. Thus, disquotability should not be considered as sufficient for excluding a term from the context-sensitivity club. Finally, I have sided with Finlay (2014) in arguing that the disagreement problem is to be addressed in a piecemeal fashion. To a large degree, our disagreement intuitions can be explained by referring to a conflict in preferences. Where speakers explicitly deny another person’s claim by stating ‘That is false’ or ‘What you said is wrong’, these speakers fall prey to a mild form of semantic blindness.

Chapter 8

Conclusion

In chapter 2, I introduced Doxastic Contextualism and located it in the current semantics literature. I showed that the Kratzer-style contextualist semantics for the information-sensitivity of the deontic ‘should’ provided by Silk (2012), Dowell (2013), and Carr (2015) can model the information-sensitivity of the doxastic ‘should’. That Charlow (2013) and Cariani’s (2013) semantics cannot do this is an interesting finding of my investigation and is evidence against their theories. Subsequently, I presented a simplified account of the truth-conditions of sentences of the form ‘*S* should have *D* to *P*’ according to which they are true at a context of utterance iff the information-state provided by this context licenses adopting *D* to *P*.

I discussed a first instance of this context-sensitivity of the doxastic ‘should’, the distinction between the subjective and the objective doxastic ‘should’, in chapter 3. Whereas it probably wouldn’t be contested that the doxastic ‘should’ has a subjective sense, it is certainly more controversial to claim that it has an objective sense, too. Hence, I provided evidence that in ordinary discourse we sometimes talk about what we objectively should believe. The upshot of this for normative epistemology is that there is no conflict between the truth norm for belief and more subjective norms, like the justification or knowledge norm.

The focus in chapter 4 was on higher-order evidence and the debate about it between defenders of Steadfastness and Conciliationism. I argued that their positions are actually not opposed to each other. We can accept both since those in favour of Steadfastness are talking about what one should believe in light of one’s

first-order evidence, whereas Conciliationists are talking about what one should believe in light of one's higher-order evidence. Accommodating both positions is desirable as there are strong arguments in support of both.

Chapter 5 and 6 turned to the practice of epistemic advice. Since this practice has so far not been discussed in the literature, I provided an account of it in chapter 5. Then, I argued for **Collectivism**—the view that in contexts where we give epistemic advice, we're employing what I called the collective doxastic 'should' rather than, for example, the subjective or objective one. The doxastic 'should' is collective when it is relativized to the collective evidence, the combined evidence of the speaker and their addressees. A nice feature of **Collectivism** is that it solves Turri's (2012) supposedly unsolvable puzzle about withholding. In chapter 6, I looked closer at how **Collectivism** could be precisified. I discussed different ways how to pool the speakers' and the addressees' evidence into the collective evidence; I expressed scepticism, for the most part, that the literature on group justification can be of much help for the purpose of precisifying **Collectivism**; and I argued that advice of the form '*S* should believe *P*' can be true even if *P* is false.

Chapter 7 closed this thesis with a defence of Doxastic Contextualism against several objections of relativists against contextualist semantics of deontic modals in general. Concerning the objection that on such semantics advisors are not addressing their advisees' questions, I showed why Dowell's (2013) news-sensitive contextualism for the practical 'should' cannot be used to solve the problem for the doxastic 'should'. Instead, I proposed that since advisors are addressing the question with the same Kaplanian character as their advisees, they can be said to address the same question, even if it does not have the same Kaplanian content. Concerning the problems surrounding speech reports and retraction, I challenged the data relativists present, and in the case of disagreement I agreed with Finlay (2014) that speakers suffer from semantic blindness if they think they disagree with others who are using another sense of 'should' than they are.

In general, I hope to have shown that the doxastic 'should' is fully information-sensitive and that this sheds light on a number of debates in normative epistemology. I believe this raises a variety of questions future research could address. One important issue is how the different senses of the doxastic 'should' relate to each other. While I have tried to stay neutral on this, I have suggested in

chapter 4 that at least some of them stand in a hierarchical relation to each other, namely, that believing what one should_{HOE} is second-best to believing what one should_{FOE}. A full account of the different senses would describe the network of their relations. This network could be strictly hierarchical (with, for example, the objective doxastic ‘should’ at the top); strictly egalitarian, such that no sense is derived from another; or a mix of these two alternatives. Another topic for future research is epistemic advice. Since it has not found much attention, there are naturally aspects of it that can be further clarified. For example, I did not try to answer whether the collective evidence is factive and how it is formed, but I only laid out options. Comparing these options could not only contribute to a better understanding of correct epistemic advice, but maybe also tell us something about the nature of evidence.

Finally, normative epistemology can benefit from a semantic analysis of the doxastic ‘should’ in other ways than by looking at its information-sensitivity. Since Kratzer-style semantics make it easy to distinguish between a doxastic and a practical ‘should’, one could conclude that debates about whether we should believe what we have epistemic reasons to believe or rather what we have practical reasons to believe are misguided and simply rooted in a confusion of language. Furthermore, in chapter 2, I mentioned that it seems to me that we rarely talk about what we must, in comparison to what we should believe. I wonder whether this could be evidence against epistemic supererogation, i.e., the phenomenon that there are propositions we should, but don’t have to believe. Lastly, Finlay (2014) proposes an end-relational semantics for ‘ought’ according to which every ought-statement is about what people should do in light of certain ends, even if they do not actually pursue those ends. His arguments, I believe, could be used to defend epistemic instrumentalism against the charge that it cannot explain why people should follow epistemic norms even if they are not pursuing an epistemic goal.¹

Long story short: semanticists should, for the foreseeable future, not run out work to do in normative epistemology.

¹I have his remarks on pragmatic catgeoricity in Finlay (2014: sec. 7.2) in mind.

Bibliography

- Adler, J. (1989). “Epistemics and the Total Evidence Requirement”, *Philosophia*, 19(2): pp. 227–243.
- Alston, W. P. (1985). “Concepts of Epistemic Justification”, *The Monist*, 68(1): pp. 57–89.
- Alston, W. P. (1988). “The Deontological Conception of Epistemic Justification”, *Philosophical Perspectives*, 2: pp. 257–299.
- Alston, W. P. (2005). *Beyond ‘Justification’: Dimensions of Epistemic Evaluation*, Ithaca: Cornell University Press.
- Bealer, G. (1998). “Intuition and the Autonomy of Philosophy”, in: M. DePaul and W. Ramsey, eds., *Rethinking Intuition: The Psychology of Intuition and Its Role in Philosophical Inquiry*, London: Rowman & Littlefield, pp. 201–240.
- Berker, S. (2013). “The Rejection of Epistemic Consequentialism”, *Philosophical Issues*, 23: pp. 363–387.
- Bird, A. (2007). “Justified Judging”, *Philosophy and Phenomenological Research*, 74(1): pp. 81–110.
- Björnsson, G. and Finlay, S. (2010). “Defending Metaethical Contextualism”, *Ethics*, 121(1): pp. 7–36.
- Blackburn, S. (1993). *Essays in Quasi-Realism*, Oxford: Oxford University Press.
- Boghossian, P. A. (2003). “The Normativity of Content”, *Philosophical Issues*, 13(1): pp. 31–45.

- Booth, A. R. (2006). “Can There Be Epistemic Reasons for Action?”, *Grazer Philosophische Studien*, 73: pp. 133–144.
- Booth, A. R. (2009). “Motivating Epistemic Reasons for Action”, *Grazer Philosophische Studien*, 78: pp. 265–271.
- Brandt, R. (1963). “Toward a Credible Form of Utilitarianism”, in: H. N. Casteñada and G. Nakhnikian, eds., *Morality and the Language of Conduct*, Detroit: Wayne State University Press, pp. 107–143.
- Brandt, R. (1967). “Epistemology and Ethics, Parallel Between”, in: P. Edwards, ed., *The Encyclopedia of Philosophy*, New York, NY: Macmillan Company and the Free Press, pp. 6–8.
- Bronfman, A. and Dowell, J. L. (forthcoming). “The Language of Reasons and Ought”, in: D. Star, ed., *Oxford Handbook of Reasons and Normativity*, Oxford: Oxford University Press.
- Broome, J. (1999). “Normative Requirements”, *Ratio*, 12(4): pp. 398–419.
- Brown, J. (2010). “Knowledge and Assertion”, *Philosophy and Phenomenological Research*, 81(3): pp. 549–566.
- Bykvist, K. and Hattiangadi, A. (2007). “Does Thought Imply Ought?”, *Analysis*, 67(296): pp. 277–285.
- Cappelen, H. and Hawthorne, J. (2009). *Relativism and Monadic Truth*, Oxford: Oxford University Press.
- Cappelen, H. and Lepore, E. (2006). “Replies”, *Philosophy and Phenomenological Research*, 78(2): pp. 469–492.
- Cariani, F. (2013). “Epistemic and Deontic Should”, *Thought: A Journal of Philosophy*, 2(1): pp. 73–84.
- Cariani, F., Kaufmann, M., and Kaufmann, S. (2013). “Deliberative Modality Under Epistemic Uncertainty”, *Linguistics and Philosophy*, 36(3): pp. 225–259.
- Carnap, R. (1950). *Logical Foundations of Probability*, Chicago: University of Chicago Press.

- Carr, J. (2015). "Subjective Ought", *Ergo*, 2(27): pp. 678–710.
- Charlow, N. (2013). "What We Know and What to Do", *Synthese*, 190(12): pp. 2291–2323.
- Charlow, N. (forthcoming). "Decision Theory: Yes! Truth Conditions: No!", in: N. Charlow and M. Chrisman, eds., *Deontic Modals*, New York, NY: Oxford University Press.
- Chrisman, M. (2006). "From Epistemic Contextualism to Epistemic Expressivism", *Philosophical Studies*, 135(2): pp. 225–254.
- Chrisman, M. (2008). "Ought to Believe", *Journal of Philosophy*, 105(7): pp. 346–370.
- Christensen, D. (2007). "Epistemology of Disagreement: The Good News", *Philosophical Review*, 116(2): pp. 187–217.
- Christensen, D. (2010). "Higher-Order Evidence", *Philosophy and Phenomenological Research*, 81(1): pp. 185–215.
- Christensen, D. (2011). "Disagreement, Question-Begging and Epistemic Self-Criticism", *Philosophers' Imprint*, 11(6): pp. 1–22.
- Christensen, D. (2013). "Epistemic Modesty Defended", in: D. Christensen and J. Lackey, eds., *The Epistemology of Disagreement: New Essays*, Oxford University Press, pp. 76–97.
- Clifford, W. K. (1877/1999). "The Ethics of Belief", in: *The Ethics of Belief and Other Essays*, Amherst, NY: Prometheus Books, pp. 70–96.
- Coates, A. (2012). "Rational Epistemic Akrasia", *American Philosophical Quarterly*, 49(2): pp. 113–124.
- Cohen, S. (1986). "Knowledge and Context", *Journal of Philosophy*, 83(10): pp. 574–583.
- Cohen, S. (1998). "Contextualist Solutions to Epistemological Problems: Scepticism, Gettier, and the Lottery", *Australasian Journal of Philosophy*, 76(2): pp. 289–306.

- Cohen, S. (ms). “Theorizing about the Epistemic”, *Unpublished Manuscript*, URL <http://www.stew-cohen.com/storage/Theorizing%20about%20the%20epistemic-final.pdf>.
- Comesaña, J. (2013). “On a Puzzle about Withholding”, *The Philosophical Quarterly*, 63(251): pp. 374–376.
- DeRose, K. (1991). “Epistemic Possibilities”, *Philosophical Review*, 100(4): pp. 581–605.
- DeRose, K. (1992). “Contextualism and Knowledge Attributions”, *Philosophy and Phenomenological Research*, 52(4): pp. 913–929.
- Dogramaci, S. (2012). “Reverse Engineering Epistemic Evaluations”, *Philosophical and Phenomenological Research*, 84(3): pp. 513–530.
- Dowell, J. L. (2011). “A Flexible Contextualist Account of Epistemic Modals”, *Philosophers’ Imprint*, 11(14): pp. 1–25.
- Dowell, J. L. (2012). “Contextualist Solutions to Three Puzzles about Practical Conditionals”, in: R. Shafer-Landau, ed., *Oxford Studies in Metaethics*, vol. 7, Oxford: Oxford University Press, pp. 271–303.
- Dowell, J. L. (2013). “Flexible Contextualism about Deontic Modals: A Puzzle about Information-Sensitivity”, *Inquiry*, 56(2-3): pp. 149–178.
- Egan, A. (2007). “Epistemic Modals, Relativism and Assertion”, *Philosophical Studies*, 133(1).
- Egan, A., Hawthorne, J., and Weatherson, B. (2005). “Epistemic Modals in Context”, in: G. Preyer and G. Peter, eds., *Contextualism in Philosophy*, Oxford: Oxford University Press, pp. 131–170.
- Elga, A. (2007). “Reflection and Disagreement”, *Nous*, 41(3): pp. 478–502.
- Engel, P. (2007). “Belief and Normativity”, *Disputatio*, 2(23): pp. 179–203.
- Engel, P. (2013). “Doxastic Correctness”, *Aristotelian Society Supplementary Volume*, 87(1): pp. 199–216.
- Ewing, A. C. (1953). *Ethics*, London: English Universities Press.

- Fantl, J. and McGrath, M. (2002). “Evidence, Pragmatics, and Justification”, *Philosophical Review*, 111(1): pp. 67–94.
- Feldman, R. (1988). “Subjective and Objective Justification in Ethics and Epistemology”, *The Monist*, 71(3): pp. 405–419.
- Feldman, R. (1995). “Authoritarian Epistemology”, *Philosophical Topics*, 23(1): pp. 147–169.
- Feldman, R. (2000). “The Ethics of Belief”, *Philosophy and Phenomenological Research*, 60(3): pp. 667–695.
- Feldman, R. (2005). “Respecting the Evidence”, *Philosophical Perspectives*, 19: pp. 95–119.
- Feldman, R. (2009). “Evidentialism, Higher-Order Evidence, and Disagreement”, *Episteme*, 6(3): pp. 294–312.
- Field, H. (2009). “Epistemology Without Metaphysics”, *Philosophical Studies*, 143(2): pp. 249–290.
- Finlay, S. (2014). *Confusion of Tongues: A Theory of Normative Language*, New York, NY: Oxford University Press.
- Finlay, S. (forthcoming). “‘Ought’: Out of Order”, in: N. Charlow and M. Chrisman, eds., *Deontic Modality*, Oxford: Oxford University Press.
- Foley, R. (1987). *The Theory of Epistemic Rationality*, Cambridge, MA: Harvard University Press.
- Fumerton, R. (1990). “Metaepistemology and Skepticism”, in: M. Roth and G. Ross, eds., *Doubting: Contemporary Perspectives on Skepticism*, Dordrecht: Kluwer, pp. 57–68.
- Gauker, C. (2008). “Zero Tolerance for Pragmatics”, *Synthese*, 165(3): pp. 359–371.
- Gibbard, A. (2003). *Thinking How to Live*, Cambridge: Harvard University Press.
- Gibbard, A. (2005). “Truth and Correct Belief”, *Philosophical Issues*, 15(1): pp. 338–350.

- Gibbons, J. (2013). *The Norm of Belief*, Oxford: Oxford University Press.
- Ginet, C. (2001). “Deciding to Believe”, in: M. Steup, ed., *Knowledge, Truth, and Duty*, Oxford: Oxford University Press, pp. 63–76.
- Glanzberg, M. (2007). “Context, Content, and Relativism”, *Philosophical Studies*, 136(1): pp. 1–29.
- Glüer, K. and Wikforss, A. (2009). “Against Content Normativity”, *Mind*, 118(469): pp. 31–70.
- Glüer, K. and Wikforss, A. (2010). “The Truth Norm and Guidance: A Reply to Steglich-Petersen”, *Mind*, 119(475): pp. 757–761.
- Glüer, K. and Wikforss, Å. (2015). “Still No Guidance: Reply to Steglich-Petersen”, *Theoria*, 81(3): pp. 272–279.
- Goldman, A. (1979). “What Is Justified Belief?”, in: G. S. Pappas, ed., *Justification and Knowledge*, Dordrecht: D. Reidel, pp. 89–104.
- Goldman, A. (1986). *Epistemology and Cognition*, Cambridge, MA: Harvard University Press.
- Goldman, A. (2010). “Why Social Epistemology Is Real Epistemology”, in: A. Haddock, A. Millar, and D. Pritchard, eds., *Social Epistemology*, Oxford: Oxford University Press, pp. 1–29.
- Goldman, A. (2014). “Social Process Reliabilism”, in: J. Lackey, ed., *Essays in Collective Epistemology*, Oxford: Oxford University Press, pp. 11–39.
- Goldman, A. and Blanchard, T. (2015). “Social Epistemology”, in: E. N. Zalta, ed., *Stanford Encyclopedia of Philosophy*, URL <http://plato.stanford.edu/entries/epistemology-social/>.
- Greaves, H. and Wallace, D. (2006). “Justifying Conditionalization: Conditionalization Maximizes Expected Epistemic Utility”, *Mind*, 115(459): pp. 607–632.
- Gross (2001). *Essays on Linguistic Context-Sensitivity and Its Philosophical Significance*, New York, NY, and London: Routledge.

- Hakli, R. (2011). “On Dialectical Justification of Group Beliefs”, in: H. B. Schmid, D. Sirtes, and M. Weber, eds., *Collective Epistemology*, Frankfurt a. M.: Ontos, pp. 119–153.
- Harman, G. (1993). “Review: Stringency of Rights and ‘Ought’”, *Philosophy and Phenomenological Research*, 53(1): pp. 181–185.
- Hawthorne, J. (2004). *Knowledge and Lotteries*, Oxford: Oxford University Press.
- Hedberg, T. (2014). “Epistemic Supererogation and Its Implications”, *Synthese*, 191(15): pp. 3621–3637.
- Heim, I. (1982). *The Semantics of Definite and Indefinite Noun Phrases*, Ph.D. thesis, University of Massachusetts, Amherst.
- Heim, I. and Kratzer, A. (1998). *Semantics in Generative Grammar*, Malden, MA: Blackwell.
- Hinchman, E. S. (2005). “Advising as Inviting to Trust”, *Canadian Journal of Philosophy*, 35(3): pp. 355–386.
- Hobbes, T. (1651/2009). *Leviathan*, Oxford: Oxford University Press.
- Horowitz, S. (2014). “Epistemic Akrasia”, *Nous*, 48(4): pp. 718–744.
- Humberstone, L. (2006). “Sufficiency and Excess”, *Aristotelian Society Supplementary Volume*, 80(1): pp. 265–320.
- Hume, D. (1748/2007). *An Enquiry Concerning Human Understanding*, Oxford: Oxford University Press.
- Huvenes, T. T. (2012). “Varieties of Disagreement and Predicates of Taste”, *Australasian Journal of Philosophy*, 90(1): pp. 167–181.
- Huvenes, T. T. (2015). “Epistemic Modals and Credal Disagreement”, *Philosophical Studies*, 172(4): pp. 987–1011.
- Jackson, F. (1986). “A Probabilistic Approach to Moral Responsibility”, in: R. Barcan Marcus, G. Dorn, and P. Weingartner, eds., *Logic, Methodology and Philosophy of Science VII*, Amsterdam: Elsevier Science Publisher, pp. 351–365.

- Jackson, F. (1991). “Decision-Theoretic Consequentialism and the Nearest and Dearest Objection”, *Ethics*, 101(3): pp. 461–482.
- Jeffrey, R. C. (1983). *The Logic of Decision*, Chicago: University of Chicago Press.
- Jenkins, C. S. (2007). “Epistemic Norms and Natural Facts”, *American Philosophical Quarterly*, 44.
- Joyce, J. M. (1998). “A Nonpragmatic Vindication of Probabilism”, *Philosophy of Science*, 65(4): pp. 575–603.
- Kamp, H. (1981). “A Theory of Truth and Semantic Representation”, in: J. Groenendijk, T. Janssen, and M. Stokhof, eds., *Formal Methods in the Study of Language*, Amsterdam: Mathematisch Centrum, pp. 277–322.
- Kaplan, D. (1978). “Dthat”, in: P. Cole, ed., *Syntax and Semantics*, New York, NY: Academic Press, pp. 221–243.
- Kaplan, D. (1989a). “Afterthoughts”, in: J. Almog, J. Perry, and H. Wettstein, eds., *Themes from Kaplan*, Oxford: Oxford University Press, pp. 565–614.
- Kaplan, D. (1989b). “Demonstratives. An Essay on the Semantics, Logic, Metaphysics and Epistemology of Demonstratives and Other Indexicals”, in: J. Almog, J. Perry, and H. Wettstein, eds., *Themes from Kaplan*, Oxford: Oxford University Press, pp. 481–563.
- Kelly, T. (2003). “Epistemic Rationality as Instrumental Rationality: A Critique”, *Philosophy and Phenomenological Research*, 66(3): pp. 612–640.
- Kelly, T. (2005). “The Epistemic Significance of Disagreement”, in: T. S. Gendler and J. Hawthorne, eds., *Oxford Studies in Epistemology*, vol. 1, Oxford: Oxford University Press, pp. 167–196.
- Kelly, T. (2007). “Evidence and Normativity: Reply to Leite”, *Philosophy and Phenomenological Research*, 75(2): pp. 465–474.
- Kelly, T. (2008). “Evidence: Fundamental Concepts and the Phenomenal Conception”, *Philosophy Compass*, 3(5): pp. 933–955.

- Kelly, T. (2010). “Peer Disagreement and Higher-Order Evidence”, in: R. Feldman and T. A. Warfield, eds., *Disagreement*, Oxford: Oxford University Press, pp. 111–172.
- Kelly, T. (2014). “Evidence”, in: E. N. Zalta, ed., *Stanford Encyclopedia of Philosophy*, URL <http://plato.stanford.edu/entries/evidence/>.
- King, J. C. (2014). “Speaker Intentions in Context”, *Nous*, 48(2): pp. 219–237.
- Kölbel, M. (2002). *Truth without Objectivity*, London: Routledge.
- Kolodny, N. and MacFarlane, J. (2010). “Ifs and Oughts”, *Journal of Philosophy*, 107(3): pp. 115–143.
- Kornblith, H. (1983). “Justified Belief and Epistemically Responsible Action”, *Philosophical Review*, 91(1): pp. 33–48.
- Kornblith, H. (1985). “Ever Since Descartes”, *The Monist*, 68(2): pp. 264–276.
- Kornblith, H. (2002). *Knowledge and its Place in Nature*, Oxford: Oxford University Press.
- Kratzer, A. (1977). “What ‘Must’ and ‘Can’ Must and Can Mean”, *Linguistics and Philosophy*, 1(3): pp. 337–355.
- Kratzer, A. (1981). “The Notional Category of Modality”, in: H.-J. Eikmeyer and H. Rieser, eds., *Words, Worlds, and Contexts: New Approaches in Word Semantics*, Berlin: de Gruyter, pp. 38–74.
- Kratzer, A. (1991). “Modality”, in: A. von Stechow and D. Wunderlich, eds., *Semantics: An International Handbook of Contemporary Research*, Berlin: de Gruyter, pp. 639–650.
- Kvanvig, J. L. (1984). “Subjective Justification”, *Mind*, 93(369): pp. 71–84.
- Kvanvig, J. L. (2003). *The Value of Knowledge and the Pursuit of Understanding*, Cambridge: Cambridge University Press.
- Kvanvig, J. L. and Menzel, C. (1990). “The Basic Notion of Justification”, *Philosophical Studies*, 59(3): pp. 235–261.

- Lackey, J. (2010). “What Should We Do When We Disagree?”, in: J. Hawthorne and Z. G. Szabo, eds., *Oxford Studies in Epistemology*, vol. 3, Oxford: Oxford University Press, pp. 274–293.
- Lackey, J. (ms). “Group Belief: Lessons from Lies and Bullshit”, *Unpublished Manuscript*.
- Lasersohn, P. (2005). “Context Dependence, Disagreement, and Predicates of Personal Taste”, *Linguistics and Philosophy*, 28(6): pp. 643–686.
- Lasonen-Aarnio, M. (2014). “Higher-Order Evidence and the Limits of Defeat”, *Philosophy and Phenomenological Research*, 88(2): pp. 314–345.
- Lasonen-Aarnio, M. (ms). “Enkrasia or Evidentialism? Learning to Love Mismatch”, *Unpublished Manuscript*.
- Lehrer, K. and Cohen, S. (1983). “Justification, Truth, and Coherence”, *Synthese*, 55(2): pp. 191–207.
- Lewis, D. (1973). *Counterfactuals*, Cambridge, MA: Harvard University Press.
- List, C. and Pettit, P. (2002). “Aggregating Sets of Judgments: An Impossibility Result”, *Economics and Philosophy*, 18(1): pp. 89–110.
- List, C. and Pettit, P. (2004). “Aggregating Sets of Judgments: Two Impossibility Results Compared”, *Synthese*, 140(1-2).
- Littlejohn, C. (2011). “Evidence and Knowledge”, *Erkenntnis*, 74(2): pp. 241–262.
- Littlejohn, C. (2012). *Justification and the Truth-Connection*, Cambridge: Cambridge University Press.
- Lord, E. (2014). “Review of *The Norm of Belief* by J. Gibbons”, *Notre Dame Philosophical Reviews*.
- MacFarlane, J. (2003). “Future Contingents and Relative Truth”, *The Philosophical Quarterly*, 53(212): pp. 321–336.
- MacFarlane, J. (2005). “The Assessment Sensitivity of Knowledge Attributions”, in: T. S. Gendler and J. Hawthorne, eds., *Oxford Studies in Epistemology*, vol. 1, Oxford: Oxford University Press, pp. 197–233.

- MacFarlane, J. (2007). “Relativism and Disagreement”, *Philosophical Studies*, 132(1): pp. 17–31.
- MacFarlane, J. (2009). “Nonindexical Contextualism”, *Synthese*, 166(2): pp. 231–250.
- MacFarlane, J. (2011). “Epistemic Modals Are Assessment-Sensitive”, in: A. Egan and B. Weatherson, eds., *Epistemic Modality*, Oxford: Oxford University Press, pp. 144–178.
- MacFarlane, J. (2014). *Assessment Sensitivity: Relative Truth and Its Applications*, Oxford: Oxford University Press.
- Mackie, J. L. (1977). *Ethics: Inventing Right and Wrong*, Harmondsworth: Penguin.
- Marra, A. (2014). “For a Dynamic Semantics of Necessity Deontic Modals”, in: F. Cariani, D. Grossi, J. Meheus, and X. Parent, eds., *Deontic Logic and Normative Systems*, Cham: Springer, pp. 124–138.
- Matheson, J. (2009). “Conciliatory Views of Disagreement and Higher-Order Evidence”, *Episteme*, 6(3): pp. 269–279.
- McElwee, B. (ms). “Supererogation Across Normative Domains”, *Unpublished Manuscript*.
- McHugh, C. (2012). “The Truth Norm of Belief”, *Pacific Philosophical Quarterly*, 93(1): pp. 8–30.
- Montmarquet, J. (1986). “The Voluntariness of Belief”, *Analysis*, 46(1): pp. 49–53.
- Moser, P. K. (1985). *Empirical Justification*, Dordrecht: D. Reidel.
- Nelson, M. T. (2010). “We Have No Positive Epistemic Duties”, *Mind*, 119(473): pp. 83–102.
- Olson, J. (2011). “Error Theory and Reasons for Belief”, in: A. Reisner and A. Steglich-Petersen, eds., *Reasons for Belief*, Oxford: Oxford University Press, pp. 75–93.

- Papineau, D. (2013). "There Are No Norms of Belief", in: T. Chan, ed., *The Aim of Belief*, Oxford: Oxford University Press, pp. 64–97.
- Parfit, D. (1997). "Reasons and Motivation", *Aristotelian Society Supplementary Volume*, 71(1): pp. 99–130.
- Parfit, D. (2011). *On What Matters*, vol. 1, Oxford: Oxford University Press.
- Parfit, D. (ms). "What We Together Do", *Unpublished Manuscript*.
- Perry, J. (2009). "Directing Intentions", in: J. Almog and P. Leonardi, eds., *The Philosophy of David Kaplan*, New York, NY: Oxford University Press, pp. 187–207.
- Pollock, J. L. (1979). "A Plethora of Epistemological Theories", in: G. S. Pappas, ed., *Justification and Knowledge*, Dordrecht: D. Reidel, pp. 93–113.
- Pollock, J. L. (2008). "Irrationality and Cognition", in: Q. Smith, ed., *Epistemology: New Essays*, Oxford: Oxford University Press, pp. 249–274.
- Pollock, J. L. and Cruz, J. (1999). *Contemporary Theories of Knowledge*, Lanham, MD: Rowman & Littlefield.
- Portner, P. (2009). *Modality*, Oxford University Press.
- Prichard, H. (1932). "Duty and Ignorance of Fact", *Proceedings of the British Academy*, 18: pp. 67–92.
- Raz, J. (2009). "Reasons: Practical and Adaptive", in: D. Sobel and S. Wall, eds., *Reasons for Action*, Cambridge: Cambridge University Press, pp. 37–57.
- Regan, D. (1980). *Utilitarianism and Co-operation*, Oxford: Oxford University Press.
- Richard, M. (2004). "Contextualism and Relativism", *Philosophical Studies*, 119(1-2): pp. 215–242.
- Ridge, M. (2007). "Epistemology for Ecumenical Expressivists", *Aristotelian Society Supplementary Volume*, 81(1): pp. 83–105.
- Ridge, M. (2014). *Impassioned Belief*, Oxford: Oxford University Press.

- Ross, W. D. (1939). *Foundations of Ethics*, Oxford: Oxford University Press.
- Schaffer, J. (2004). “From Contextualism to Contrastivism”, *Philosophical Studies*, 119(1-2): pp. 73–103.
- Schmitt, F. (1994). “The Justification of Group Beliefs”, in: F. Schmitt, ed., *Socializing Epistemology: The Social Dimensions of Knowledge*, London: Rowman & Littlefield, pp. 257–287.
- Schoenfield, M. (2012). “Chilling Out on Epistemic Rationality”, *Philosophical Studies*, 158(2): pp. 197–219.
- Schoenfield, M. (2015). “A Dilemma for Calibrationism”, *Philosophy and Phenomenological Research*, 91(2): pp. 425–455.
- Schoenfield, M. (ms). “An Accuracy Based Approach to Higher Order Evidence”, *Unpublished Manuscript*, URL <http://www.miriamschoenfield.com/research/accuracy-based-approach.pdf>.
- Schroeder, M. (2008). *Being For*, Oxford: Oxford University Press.
- Schroeder, M. (2015). “Knowledge Is Belief For Sufficient (Objective and Subjective) Reason”, in: T. S. Gendler and J. Hawthorne, eds., *Oxford Studies in Epistemology*, vol. 5, Oxford: Oxford University Press, pp. 226–252.
- Searle, J. R. (1969). *Speech Acts: An Essay in the Philosophy of Language*, Cambridge: Cambridge University Press.
- Sennet, A. (2016). “Ambiguity”, in: E. N. Zalta, ed., *Stanford Encyclopedia of Philosophy*, URL <http://plato.stanford.edu/entries/ambiguity/>.
- Shah, N. (2003). “How Truth Governs Belief”, *Philosophical Review*, 112(4): pp. 447–482.
- Shah, N. and Velleman, J. D. (2005). “Doxastic Deliberation”, *Philosophical Review*, 114(4): pp. 497–534.
- Shope, R. (1978). “The Conditional Fallacy in Contemporary Philosophy”, *Journal of Philosophy*, 75(8): pp. 397–413.
- Sidgwick, Henry (1874). *The Methods of Ethics*, London: Macmillan.

- Silk, A. (2012). “Modality, Weights, and Inconsistent Premise Sets”, in: A. Chereches, ed., *Proceedings of SALT 22*, Ithaca, NY: CLC Publications.
- Silk, A. (2014). “Evidence Sensitivity in Weak Necessity Deontic Modals”, *Journal of Philosophical Logic*, 43(4): pp. 691–723.
- Sloman, A. (1970). “‘Ought’ and ‘Better’”, *Mind*, 79(315): pp. 385–394.
- Smithies, D. (2012). “The Normative Role of Knowledge”, *Nous*, 46(2): pp. 265–288.
- Sosa, E. (2001). “For The Love of Truth?”, in: L. Zagzebski and A. Fairweather, eds., *Virtue Epistemology: Essays on Epistemic Virtue and Responsibility*, Oxford: Oxford University Press, pp. 49–62.
- Stalnaker, R. (1997). “Reference and Necessity”, in: B. Hale and C. Wright, eds., *A Companion to the Philosophy of Language*, Oxford: Blackwell, pp. 534–554.
- Stanley, J. (2005). *Knowledge and Practical Interests*, Oxford: Oxford University Press.
- Stanley, J. and Gendler Szabó, Z. (2000). “On Quantifier Domain Restriction”, *Mind & Language*, 15(2-3): pp. 219–261.
- Starr, W. (forthcoming). “Dynamic Expressivism about Deontic Modality”, in: N. Charlow and M. Chrisman, eds., *Deontic Modality*, Oxford: Oxford University Press.
- Steup, M. (2000). “Doxastic Voluntarism and Epistemic Deontology”, *Acta Analytica*, 15(1): pp. 25–56.
- Stevenson, C. (1944). *Ethics and Language*, New Haven, CT: Yale University Press.
- Stokke, A. (2010). “Intention-Sensitive Semantics”, *Synthese*, 175(3): pp. 383–404.
- Sutton, J. (2007). *Without Justification*, Cambridge, MA: MIT Press.
- Thomson, J. J. (1986). “Imposing Risks”, in: W. Parent, ed., *Rights, Restitution, and Risk: Essays in Moral Theory*, Cambridge, MA: Harvard University Press, pp. 173–191.

- Turri, J. (2012). “A Puzzle about Withholding”, *The Philosophical Quarterly*, 62(247): pp. 355–364.
- Unger, P. (1984/2002). *Philosophical Relativity*, Oxford: Oxford University Press.
- Van Inwagen, P. (1996). “It Is Wrong, Everywhere, Always, and for Anyone to Believe Anything on Insufficient Evidence”, in: J. Jordan and D. Howard-Snyder, eds., *Faith, Freedom, and Rationality*, Lanham, MA: Rowman and Littlefield, pp. 137–154.
- von Fintel, K. and Gillies, A. S. (2008). “CIA Leaks”, *Philosophical Review*, 117(1): pp. 77–98.
- von Fintel, K. and Heim, I. (2011). “Intensional Semantics”, *Lecture Notes, MIT*.
- von Fintel, K. and Iatridou, S. (2008). “How to Say *Ought* in Foreign: the Composition of Weak Necessity Modals”, in: J. Guéron and J. Lecarme, eds., *Time and Modality*, Dordrecht: Springer, pp. 115–141.
- von Fintel, K. and Iatridou, S. (ms). “What to Do if You Want to Go to Harlem: Anankastic Conditionals and Related Matters”, *Unpublished Manuscript*, URL <http://web.mit.edu/fintel/www/harlem-rutgers.pdf>.
- Weatherson, B. (2003). “What Good Are Counterexamples?”, *Philosophical Studies*, 115(1): pp. 1–31.
- Weatherson, B. (ms). “Do Judgments Screen Evidence?”, *Unpublished Manuscript*, URL <http://brian.weatherson.org/JSE.pdf>.
- Wedgwood, R. (2002). “The Aim of Belief”, *Philosophical Perspectives*, 16: pp. 267–297.
- Wedgwood, R. (2007). *The Nature of Normativity*, Oxford: Oxford University Press.
- Wedgwood, R. (2013). “The Right Thing to Believe”, in: T. Chan, ed., *The Aim of Belief*, Oxford: Oxford University Press, pp. 123–139.
- Wedgwood, R. (forthcoming). “Objective and Subjective ‘Ought’”, in: N. Charlow and M. Chrisman, eds., *Deontic Modality*, Oxford: Oxford University Press.

- Weisberg, J. (2011). “Varieties of Bayesianism”, in: D. M. Gabbay, S. Hartmann, and J. Woods, eds., *Handbook of the History of Logic: Inductive Logic*, vol. 1, Amsterdam: North-Holland, pp. 477–552.
- Whitcomb, D. (2014). “Can There Be a Knowledge-First Ethics of Belief?”, in: J. Matheson and R. Vitz, eds., *The Ethics of Belief: Individual and Social*, Oxford: Oxford University Press, pp. 89–111.
- Whiting, D. (2014). “Reasons for Belief, Reasons for Action, the Aim of Belief, and the Aim of Action”, in: C. Littlejohn and J. Turri, eds., *Epistemic Norms*, Oxford: Oxford University Press, pp. 219–238.
- Wiland, E. (2000). “Advice and Moral Objectivity”, *Philosophical Papers*, 29(1): pp. 1–19.
- Willer, M. (forthcoming). “Dynamic Foundations for Deontic Logic”, in: N. Charlow and M. Chrisman, eds., *Deontic Modality*, Oxford: Oxford University Press.
- Williams, B. (1981). “Practical Necessity”, in: B. Williams, ed., *Moral Luck: Philosophical Papers 1973-1980*, Oxford: Oxford University Press, pp. 124–131.
- Williamson, J. (2009). “Aggregating Judgements by Merging Evidence”, *Journal of Logic and Computation*, 19(3): pp. 461–473.
- Williamson, J. (2015). “Deliberation, Judgement and the Nature of Evidence”, *Economics and Philosophy*, 31(1): pp. 27–65.
- Williamson, T. (2000). *Knowledge and Its Limits*, Oxford: Oxford University Press.
- Williamson, T. (forthcoming). “Justifications, Excuses, and Sceptical Scenarios”, in: F. Dorsch and J. Dutant, eds., *The New Evil Demon Problem*, Oxford: Oxford University Press.
- Worsnip, A. (2015a). “Moral Reasons, Epistemic Reasons, and Rationality”, *The Philosophical Quarterly*, 66(263): pp. 341–361.
- Worsnip, A. (2015b). “The Conflict of Evidence and Coherence”, *Philosophy and Phenomenological Research*, Online First.

Yalcin, S. (2011). “Nonfactualism About Epistemic Modality”, in: A. Egan and B. Weatherson, eds., *Epistemic Modality*, Oxford: Oxford University Press, pp. 295–332.

Yalcin, S. (2012). “Dynamic Semantics”, in: D. G. Fara and G. Russell, eds., *Routledge Companion to the Philosophy of Language*, New York, NY and London: Routledge, pp. 253–279.