

REDUCTION IN EPISTEMOLOGY

by

ROBERT BEDDOR

A dissertation submitted to the

Graduate School-New Brunswick

Rutgers, The State University of New Jersey

In partial fulfillment of the requirements

For the degree of

Doctor of Philosophy

Graduate Program in Philosophy

Written under the direction of

Andy Egan and Alvin Goldman

And approved by

New Brunswick, New Jersey

May, 2016

ABSTRACT OF THE DISSERTATION

Reduction in Epistemology

by BOB BEDDOR

Dissertation Directors:

Andy Egan and Alvin Goldman

This dissertation centers on two questions: (1) Can we explain epistemic facts in terms of non-epistemic facts? (2) What is the most explanatorily basic notion within the epistemic realm?

Many philosophers are attracted to the idea that the epistemic is reducible to the natural: facts about epistemic justification, knowledge, and the like can be explained in terms of non-epistemic facts. How could such a reduction be achieved? Chapter One explores the two leading proposals in the literature: process reliabilism and mentalist evidentialism. I argue that both of these approaches flounder when it comes to explaining epistemic defeat (cases where an individual gets some evidence in favor of a belief, which is then trumped by countervailing evidence). The standard process reliabilist treatment of defeat faces counterexamples, and leading evidentialist treatments of defeat either fall victim to the same fate or fail to be reductive.

Chapter Two explores a different reductive strategy. I suggest that we should pursue the path of semantic ascent: we should focus on epistemic linguistic expressions and seek to define them without recourse to epistemic notions. Specifically, I develop an ‘attitudinal’ semantics for a variety of epistemic expressions, according to which epistemic expressions are analyzed in terms of the conative attitudes that give rise to them. The resulting semantics is reductive; in addition, it offers to explain some of the striking commonalities between epistemic and ethical discourse.

Chapter Three considers reduction within the epistemic domain. Recent epistemology

has seen the rise of ‘Knowledge First’ epistemology, according to which knowledge is the most basic explanatorily basic epistemic notion. I advance an alternative picture, according to which *epistemic certainty* is explanatorily fundamental. After developing a context-sensitive semantics for certainty ascriptions, I go on to put certainty to explanatory work. I argue that a wide range of epistemic phenomena—including epistemic modals, evidential probability, and knowledge itself—can be profitably analyzed in terms of certainty.

Acknowledgements

I am extremely grateful to all of the people who have supported me on this long journey.

As an undergraduate at Brown I had the good fortune to have two outstanding thesis advisors, Charles Larmore and Josh Schechter, who were extremely generous with their time and comments. Over the course of long meetings in their offices they coaxed my inchoate ideas on skepticism into a more-or-less focused undergraduate thesis. More importantly, they helped me see how exciting—and how challenging—it can be to do philosophy at a serious level.

At Rutgers I have learned from some truly fantastic philosophers. My greatest debt is to my dissertation directors, Andy Egan and Alvin Goldman. During my graduate career I developed a rather wide range of interests, spanning epistemology, semantics, metaethics, and metaphysics. Alvin and Andy were unfailingly supportive of all of my wide-ranging and occasionally idiosyncratic projects. Conversations with them were invariably helpful, regardless of whether I was thinking through an idea for the first time or showing them the latest draft of a paper that they had commented on umpteen times before. Their probing questions and thoughtful suggestions have shaped this dissertation—and my approach to philosophy—in more ways than I can acknowledge.

I also owe heartfelt thanks to my other committee members. Without exception, Jonathan Schaffer's comments cut to the heart of the issues, which forced me to clarify my thoughts on every chapter of this dissertation. Throughout my graduate career, Ernie Sosa has been extraordinarily supportive. He was kind enough to invite me to join his dissertation group; the feedback from Ernie and the other members proved extremely helpful in developing the ideas in these pages. Chapter three of this dissertation was influenced

in many ways by Brian Weatherson's work, and I am thrilled that he agreed to come on board as my external committee member. I consider myself tremendously lucky to have had committee members who are both incredibly good philosophers and very generous human beings.

I also received a great deal of help and support from faculty who were not on my dissertation committee, in ways that often went well above and beyond the call of the duty. Special thanks to Holly Smith, who gave me detailed, insightful comments on both my teaching and chapter three. I also want to thank Branden Fitelson, for inviting me to present chapter three in his reading group, and Ernie Lepore, for all his professional advice and support.

During graduate school I probably learned the most from my fellow graduate students. When I first arrived at Rutgers I was amazed to find myself surrounded by so many smart people who enjoyed talking philosophy. That sense of amazement has yet to wear off. Among the students whose conversations I found particularly helpful, I'd like to thank David Black, Pavel Davydov, Marco Dees, Tom Donaldson, Mike Hicks, Nico Kirk-Giannini, Stephanie Leary, Mark Lee, Lisa Mirracchi, Pamela Robinson, Mary Salvaggio, and Tobias Wilsch. Special thanks go to Carlotta Pavese and Simon Goldstein. My long discussions with Carlotta helped form many of my views and philosophical interests. Simon's conversations—and extensive comments on paper drafts—helped me clarify my thoughts on crucial issues and to appreciate many points that I had overlooked.

Many thanks to my parents, who never questioned my less-than-lucrative decision to go to grad school in philosophy, and who were always there with love, support, and proofreading services over the years. And, finally, special thanks go to Laura Willson. Laura's sense of humor, her willingness to discuss ideas with me, and, most importantly, her unwavering love and support have made the years I've spent writing this dissertation truly happy.

Contents

Abstract	ii
Acknowledgements	iv
Introduction	viii
1 Reductive Accounts of Justification	1
1.1 Introduction	1
1.2 Why Process Reliabilists Need a Story about Defeat	3
1.3 The Alternative Reliable Process Account of Defeat	6
1.4 A Counterexample to Sufficiency	7
1.5 A Counterexample to Necessity	13
1.6 Looking Forward	16
1.7 Appealing to Evidence	17
1.8 Three Attempts to Reduce Evidence Possession	20
1.9 Conclusion	22
2 Towards a Reductive Metaepistemology	24
2.1 Introduction	24
2.2 Epistemic Approval	26
2.3 An Attitudinal Semantics	31
2.3.1 An Attitudinal Semantics for Deontic Modals	32
2.3.2 An Attitudinal Semantics for Justification Ascriptions	42
2.3.3 Beyond Modals and Justification Ascriptions	50

2.4	Advantages of an Attitudinal Metaepistemology	51
2.4.1	Value-Ladenness	51
2.4.2	Connection with Motivation	52
2.5	Comparison with a First-Order Reduction	59
3	The Centrality of Certainty	63
3.1	Introduction	63
3.2	Certainty: A Schematic Account	65
3.2.1	Subjective vs. Epistemic Certainty	65
3.2.2	In Favor of Scarce	68
3.2.3	Against Scarce	70
3.2.4	The Best of Both Worlds: Max Adjectives Without Scarce	70
3.2.5	Puzzling Pos	72
3.2.6	Taking Stock	76
3.3	Evidential Probability	77
3.3.1	Evidential Probabilities and Degrees of Certainty	78
3.3.2	Epistemological Advantages	81
3.3.3	Further Issues	84
3.3.4	Objections	85
3.4	Epistemic Modals	88
3.4.1	Two Analyses of Epistemic Modals	88
3.4.2	In Favor of the Certainty Analysis	89
3.4.3	An Objection to the Certainty Analysis	93
3.4.4	Taking Stock	98
3.5	Knowledge	98
3.6	Certainty and Action	100
3.7	Conclusion	105

Introduction

This dissertation consists of three papers on reduction in epistemology. The first two papers explore whether it's possible to give reductive analyses of epistemic notions such as justification, evidence, and knowledge. Chapter One argues that traditional attempts to give a reductive account of justification face hurdles when it comes to cases of defeat. I start by focusing on the most prominent reductive account of justification in the literature: process reliabilism. I argue that in order to avoid counterexamples, process reliabilists need to invoke a “No Defeaters” clause; in order to qualify as reductive, they need to go on to cash out defeat in non-epistemic terms. After raising counterexamples to the standard process reliabilist account of defeat (the Alternative Reliable Process account), I argue that the problem generalizes beyond process reliabilism: all would-be reductive accounts of justification need to accommodate cases of defeat—a task that turns out to be surprisingly difficult.¹

Chapter Two suggests tackling the problem of reduction from a different angle. Perhaps we should focus on the linguistic expressions that we use to make epistemic evaluations (e.g., terms like *justified*, *knows*, etc.), and try to analyze the meanings of these expressions in non-epistemic terms. What would such a reductive semantics look like? In the metaethics literature, there is a long tradition of explaining the meanings of ethical expressions in terms of agents' conative attitudes. Chapter Two explores how to extend this strategy to the epistemic domain.

I start by focusing on epistemic evaluations involving modals (e.g., *ought*, *should*, *must*, *may*). The leading approach to modals in the semantics literature (due to Kratzer) takes

¹Chapter One contains material from Beddor 2015a,b.

modals to quantify over the ‘best’ of a set of conversationally relevant worlds (the *modal base*). However, Kratzer’s semantics never tells us what makes a certain set of worlds the ‘best’—she leaves this job to the conversational context. On my proposal, we can use certain agents’ attitudes to determine what count as the ‘best’ worlds. Specifically, I suggest that there is a particular desire-like attitude—‘epistemic approval’—involved in epistemic evaluations; this attitude provides the relevant ranking over worlds when it comes to epistemic evaluations involving modals. Thus a sentence such as: ‘Poirot ought to believe the butler did it’ will be true iff Poirot believes the butler did it in all of the worlds in the modal base ranked highest by certain agents’ states of epistemic approval. I go on to show how this approach can be extended to other fragments of epistemic discourse, such as justification ascriptions.² The resulting analysis is reductive: epistemic evaluations are analyzed in terms of epistemic approval, which can be characterized without recourse to epistemic notions. What’s more, it fits naturally with a parallel analysis of moral evaluations, according to which claims such as ‘You (morally) ought to save the drowning child’ are analyzed in terms of a ranking provided by states of moral approval. I go on to show how this analysis sheds light on some of the commonalities between moral and epistemic evaluations—for example, the fact that both are connected with motivation.

Chapter Three shifts gears and explores reduction within the epistemic domain. Specifically, it explores the question: which epistemic notion is explanatorily fundamental (if any)? According to my proposal, the answer is: *certainty*. While the notion of epistemic certainty played a central role in the epistemology of the medieval and early modern periods, recent epistemology has tended to dismiss certainty as an unattainable ideal, focusing on knowledge instead. I argue that this is a mistake. Building on recent work on the semantics of gradable expressions, I develop a context-sensitive semantics for *certain*, according to which much of our everyday knowledge qualifies—in suitable contexts—as certain. Next, I argue that if we take certainty as our primitive notion in the epistemic domain, we can get considerable explanatory mileage out of it. A wide range of epistemic phenomena—evidential probability, epistemic modals, knowledge, and the normative constraints on credence, action and assertion—can be profitably analyzed in terms of certainty.

²The treatment of justification ascriptions contains material from Beddor forthcoming.

These three papers are to some extent independent of one another. The reader could endorse my criticisms of traditional reductive analyses in Chapter One without buying into the reductive metaepistemology developed in Chapter Two. Similarly, the reader could endorse my views on the centrality of certainty in Chapter Three while eschewing the reductive project pursued in the first two chapters. That said, the papers taken in conjunction provide a consistent, unified framework for thinking of reduction in epistemology. On the picture that emerges, a variety of epistemic notions are analyzed in terms of epistemic certainty, and certainty ascriptions (together with other epistemic expressions) are analyzed in terms of conative attitudes (specifically, states of epistemic approval).

Chapter 1

Reductive Accounts of Justification

1.1 Introduction

Epistemic facts supervene on the non-epistemic facts. Suppose one holds fixed all the facts about how an agent S formed a belief B at time t at a world w . And suppose one also holds fixed all the further facts about S 's mental states at t , together with B 's truth-value at w (as well, perhaps, as B 's truth-value at nearby worlds). Then it seems plausible that one will have thereby settled a supervenience base for all of the facts about the epistemic status of B —e.g., whether it's epistemically justified, whether it counts as knowledge, etc.

Given this, it's natural to endorse the following idea:

Reducibility: Epistemic facts are—in some sense—reducible to the non-epistemic facts.¹

Suppose we accept Reducibility. This in turn prompts the question: “How exactly does the reduction go? How do facts about justification, evidence, and knowledge reduce to non-epistemic facts?” It's natural to think that one of our jobs as epistemologists is to answer this question. In particular, it's natural to hope that when we as epistemologists are offering an analysis of a given epistemic property or state x (e.g., justification, evidence, knowledge),

¹For present purposes, I won't take a stand on exactly what sense of reducibility is at issue. However, one natural option is to spell out the reducibility in question in terms of *grounding*, which is taken to be a form of metaphysical determination corresponding to the “in virtue of” relation. According to this proposal, Reducibility amounts to the idea that all epistemic facts are fully grounded in non-epistemic facts. (For important work on grounding, see Fine 2001, 2012; Schaffer 2009; Rosen 2010. For useful overviews, see Clark and Higgins 2012; Trogdon 2013.)

the analysis in question will explain how the x -facts can be reduced to non-epistemic facts.

Thus Reducibility motivates the goal of a *reductive epistemology*. Following tradition, let us suppose our epistemological analyses take the form of universally quantified biconditionals. Thus an analysis of justification will be of the form:

Necessarily, for any subject S and any proposition p :

S is justified in believing p iff conditions $C_1...C_n$ obtain.²

Call an analysis of some epistemic state or property *reductive* just in case the right-hand side of the biconditional can be specified in non-epistemic terms.

A number of epistemologists have sought to provide reductive analyses along these lines. This reductive ambition is perhaps clearest in the writings of process reliabilists. For example, Goldman (1979) writes: “I want a theory of justified belief to specify in non-epistemic terms when a belief is justified” (90). The theory that Goldman goes on to develop is tailored to make good on this reductive ambition: facts about justification are understood in terms of facts about the reliability of an agent’s belief-forming processes, where “reliability” is taken to be a non-epistemic notion.³

While reliabilists have been particularly explicit about their reductive ambitions, a number of other epistemological views seem to owe their appeal, at least in part, to their reductive potential. Consider, for example, certain versions of mentalist internalism, according to which epistemic justification and rationality can be explained in terms of an agent’s non-factive mental states (e.g., experiences, seemings, and beliefs).⁴ Or consider various truth-tracking theories of knowledge, which seek to explain knowledge in terms of modal properties such as sensitivity, adherence, and safety.⁵ Arguably, one appealing feature of such views is that they offer to show how epistemic properties and states reduce to non-epistemic properties and states.

²Plausibly, there will be some constraints on what counts as an adequate analysis. Presumably the conditions on the right-hand side need to be finite, and not too horribly disjunctive. More controversially, one might insist that any such biconditional should be really regarded as shorthand for a thesis about the metaphysical grounds of justification ascriptions (Beddor 2015a). For present purposes, I will set such complications aside.

³In the subsequent literature, many have taken this ambition to be central to the spirit of reliabilism (see e.g., Conee and Feldman 1998: 4-5).

⁴See, for example, Lewis 1996; Conee and Feldman 2004, 2008.

⁵See, for example, Nozick 1981; Sosa 1999, 2004; Pritchard 2005; Lasonen-Aarnio 2010.

While I certainly agree that it would be nice to give reductive accounts of various epistemic phenomena, I think that providing reductive accounts is much more difficult than many have thought. In this chapter, I raise some general obstacles for reductive accounts in epistemology. For concreteness, I will focus on attempts to give a reductive account of epistemic justification. I argue that extant attempts to give a reductive account of justification have difficulty handling cases of *defeat* (cases in which an agent is *prima facie* justified in believing some proposition *p*, but this *prima facie* justification is swamped by countervailing evidence). While I take justification as my focus, I argue that the difficulties I raise also apply to reductive accounts of other epistemic states. In particular, they create difficulties for attempts to give a reductive analysis of any epistemic state that is incompatible with defeat—for example, knowledge or evidence possession.

This chapter is structured as follows. I start with process reliabilist accounts of justification, since process reliabilists have been particularly vocal about their reductive goals. I first show why process reliabilists need to supplement their view with a ‘No Defeaters’ clause (§1.2); I then argue, by way of counterexample, that the standard process reliabilist account of defeat (the ‘Alternative Reliable Process’ account) fails to articulate either necessary or sufficient conditions on defeat (§§1.3-1.5). I go on to generalize my argument: every extensionally adequate account of justification needs to deliver the right results in cases of defeat, and it is far from clear how to develop a reductive account that does so (§§1.6-1.8). I conclude by discussing the upshot for a reductive epistemology more generally (§1.9).

1.2 Why Process Reliabilists Need a Story about Defeat

To see why process reliabilists need a story about defeat, consider a simple version of process reliabilism that lacks any “No Defeaters” clause:

Simple Process Reliabilism: S’s belief that *p* is justified at *t* iff S’s belief that *p* at *t* is the result of a reliable belief-forming process (or reliable belief-forming processes).⁶

⁶Despite the amount of press it’s received, it’s not clear whether anyone has endorsed Simple Process Reliabilism, so it may be something of a strawperson. Nonetheless, it provides a useful entry point for exploring process reliabilists’ difficulties with defeat.

To see why such an account proves inadequate, consider a stock example of defeat:

Consuela and the Vase: Consuela sees a red vase in good lighting conditions at t_1 . Consequently, she comes to believe RED: $\langle \text{There's a red vase in front of me} \rangle$. At t_2 , a usually reliable informant tells Consuela that she's actually looking at a white vase illuminated by a red light. Consuela has no reason to distrust this informant; nonetheless, she disregards his testimony and continues to believe RED.

Fill in the details in the right way, and most people have the intuition that Consuela's belief in RED is defeated at t_2 : though it may be *prima facie* justified, it's not *ultima facie* justified (that is, it's not justified *full stop*). But Simple Process Reliabilism predicts precisely the opposite. After all, Consuela's belief in RED at t_2 is the result of vision operating in good lighting conditions—a paradigmatic example of a reliable cognitive process.

Is there any way to defend Simple Process Reliabilism from this objection? One might be inclined to reply as follows:

“Everyone knows that process reliabilists face the Generality Problem: they face the notoriously difficult problem of typing belief-forming processes.⁷ You've picked a fairly coarse-grained way of typing Consuela's belief-forming process at t_2 : you've characterized it as vision (or vision operating in good lighting conditions). But perhaps we should appeal to a more fine-grained characterisation of the process responsible for Consuela's belief. Let T be a function that takes as input visual experiences *together with testimony that those experiences are misleading* and produces as output belief in the content of those experiences. If we take T to be the correct way of typing Consuela's belief-forming process, it's not clear that her belief in RED was formed by a reliable process.”

Call this the ‘Typing Defense.’ In evaluating the Typing Defense, we should first note that Consuela's case is just one among many. Indeed, we can concoct an infinite variety of cases with the following structure: *S forms an intuitively justified belief that p at time t_1 . Then, at t_2 , S acquires a good reason to abandon her belief that p. Nonetheless, S sticks to her guns and continues to believe p anyway.* And so the proponent of the Typing Defense has her work cut out for her: for each such case, her proposed method for typing

⁷For a seminal statement of the Generality Problem, see Conee and Feldman 1998.

belief-forming processes will need to deliver the result that S’s belief-forming process at t_1 is reliable and that S’s belief-forming process at t_2 is unreliable.

If we turn to some of the leading attempts to solve the Generality Problem in the literature, we find that they don’t deliver these results. For example, one popular approach to the Generality Problem is to type belief-forming processes in accordance with the “common sense” classifications that the folk employ in ordinary life.⁸ Any such constraint on typing belief-forming processes bodes ill for the Typing defense. After all, the folk are far more likely to classify belief-forming processes as instances of *vision* or *vision operating in good lighting conditions* than they are to classify belief-forming processes as instances of *vision together with testimony that those visual experiences are misleading*.⁹

Of course, one might respond by simply rejecting such approaches: one might insist that the only adequate solution to the Generality Problem is a solution that’s consistent with the Typing Defense. But I see no reason to impose such a stringent adequacy condition on solutions to the Generality Problem. Typing belief-forming processes is difficult enough without worrying about defeat. Thus it would be nice if process reliabilists could give an independent treatment of defeat—a treatment of defeat that’s compatible with a variety of solutions to the Generality Problem.

Most process reliabilists agree with me on this front. Most process reliabilists reject Simple Process Reliabilism, and many of them do so—at least in part—because of Simple Process Reliabilism’s acknowledged difficulties handling cases of defeat.¹⁰ Instead of Simple Process Reliabilism, most process reliabilists opt for what we can call a ‘Two Step’ version of process reliabilism.

By a ‘Two Step’ theory of justification, I mean any theory that comprises two separate components: a *prima facie* justification condition and a ‘No Defeaters’ condition. Process

⁸Some of Goldman’s remarks in his (1979) paper suggest an approach along these lines. Jönsson (2013a,b) and Olsson (forthcoming) both defend versions of this approach.

⁹Another attractive approach to the Generality Problem is to type belief-forming processes causally. (See e.g. Goldman 1986; Becker 2008.) A natural way of developing a causal approach to the Generality Problem is to insist that if a feature f doesn’t causally affect whether S believes p at t , our way of typing the belief-forming process responsible for S’s belief that p at t shouldn’t mention f . Any such solution also stands in tension with the Typing Defense, since there’s no guarantee that defeaters will always causally affect the target belief. For example, in the case of Consuela and the Vase, we can imagine Consuela (unjustifiably) regards her interlocutor as completely unreliable; hence his testimony doesn’t causally affect her credence in RED.

¹⁰See e.g. Goldman 1979; Lyons 2009, forthcoming; Grundmann 2009; Bedke 2010.

reliabilists who pursue a Two Step theory typically take Simple Process Reliabilism (or some close variant) and convert it into an account of *prima facie* justification. For instance, the following is a straightforward version of Two Step Process Reliabilism:

S's belief that p is (*ultima facie*) justified at t iff

1. S's belief that p is the result of a reliable belief-forming process (or belief-forming processes) at t
2. S's belief that p isn't defeated at t .

A Two Step process reliabilist will allow that Consuela's belief in RED is *prima facie* justified at t_2 , but deny that it's *ultima facie* justified at t_2 , since it doesn't satisfy condition (2).¹¹

However, process reliabilists who go down this road face an obvious challenge. *Defeat* is an epistemic notion *par excellence*. Thus, on pain of abandoning their reductive aspirations, Two Step process reliabilists need to provide a reductive account of defeat.

1.3 The Alternative Reliable Process Account of Defeat

To their credit, process reliabilists have acknowledged this obligation and tried to discharge it. The classic process reliabilist story about defeat was first proposed by Goldman (1979), and has found a recent champion in Lyons (2009, forthcoming). It goes like this:

Alternative Reliable Process Account of Defeat (ARP): S's belief that p is defeated at t iff there's some reliable or conditionally reliable process available to S at t which, if it had been used by S in addition to the process actually used, would have resulted in S's not believing p at t .¹²

¹¹In addition to its superior handling of Consuela and the Vase, Two Step Process Reliabilism may fare better than Simple Process Reliabilism when it comes to dealing with Bonjour's case of Norman the clairvoyant (Bonjour 1985) and Lehrer's case of Truetemp (Lehrer 1990). As Goldman (1986) suggests, one option is to say that these characters' beliefs are *prima facie* justified but defeated. (Though see Lyons (2009) for a version of process reliabilism that precludes Norman and Truetemp's beliefs from possessing *prima facie* justification.)

¹²Grundmann (2009) and Bedke (2010) defend similar counterfactual accounts of defeat. While their accounts (particularly Grundmann's) differ from ARP in some important details, I'll forego a discussion of such details since—differences notwithstanding—the counterexamples I offer to ARP straightforwardly generalize to their proposals.

How does ARP handle Consuela and the Vase? I take it that proponents of ARP will offer the following diagnosis: Consuela's belief in RED is defeated at t_2 because she could have given weight to her interlocutor's testimony. More precisely: there's a possible process that takes his testimony as input and outputs a fairly high credence in the content of that testimony. Plausibly, this process is either reliable or conditionally reliable (or at least there's a natural way of typing this process that makes it either reliable or conditionally reliable). Consuela could have used this process in addition to visual perception; had she done so, she would have abandoned her belief in RED.

I'm happy to grant that ARP is capable of explaining our intuition about Consuela and the Vase. Unfortunately, there are at least some cases where ARP clearly delivers the wrong results.

For ease of exposition, it will be useful to explicitly distinguish ARP's sufficiency condition for defeat from ARP's necessity condition:

Sufficiency: If there's some reliable or conditionally reliable process available to S at t which, if it had been used by S in addition to the process actually used, would have resulted in S's not believing p at t , then S's belief that p is defeated at t .

Necessity: If S's belief that p is defeated at t , there's some reliable or conditionally reliable process available to S at t which, if it had been used by S in addition to the process actually used, would have resulted in S's not believing p at t .

In what follows, I first present a counterexample to Sufficiency (§1.4); next, I offer a counterexample to Necessity (§1.5).

1.4 A Counterexample to Sufficiency

Before presenting what I take to be a convincing counterexample to Sufficiency, I'll briefly mention an unconvincing counterexample. Kvanvig (2007) asks us to imagine that "[T]here is a competent cogniser who disagrees with you about something you know to be true... There is a reliable process which if you had used it would have resulted in a different belief: namely, ask this cogniser and believe what is reported." (2007: 1-2)

The reason I take Kvanvig’s counterexample to be unconvincing is that Goldman’s initial (1979) discussion of ARP makes it clear he doesn’t intend the alternative processes in question to include consulting new interlocutors. Goldman writes:

[I]t seems implausible to say all ‘available’ processes ought to be used, at least if we include such processes as gathering new evidence. Surely a belief can sometimes be justified even if additional evidence-gathering would yield a different doxastic attitude. What I think we should have in mind here are such additional process as calling previously acquired evidence to mind, assessing the implications of that evidence, etc. (Goldman 1979)

It seems clear that Goldman intends ARP to be restricted to alternative reliable cognitive processes that are in an important sense internal; they do not involve further research. And so the process suggested by Kvanvig—*consulting competent cognizers who disagree with you*—won’t qualify as ‘available’ in the relevant sense.¹³

Nonetheless, I think there are other cases with a similar structure that do present compelling counterexamples to Sufficiency. Here’s one:

Thinking About Unger: Harry sees a tree in front of him at t . Consequently, he comes to believe the proposition TREE: \langle There is a tree in front of me \rangle at t . Now, Harry happens to be very good at forming beliefs about what Peter Unger’s 1975 time-slice would advise one to believe in any situation. Call this cognitive process his “Unger Predictor”: for any proposition p , any agent A, and any situation s , Harry’s Unger Predictor spits out an accurate belief about what doxastic attitude Unger’s 1975 time-slice would advise A to take towards p in s .

¹³Given this restriction of what counts as an alternative reliable process ‘available’ to the agent, it’s not clear that Goldman is entitled to his later verdict on Kornblith’s (1983) case of Jones, “a headstrong young physicist” who gives a talk wherein he announces his belief that p . Jones is unable to withstand any sort of criticism, and consequently doesn’t listen to his colleague’s devastating objection. (It’s crucial to Kornblith’s case that Jones doesn’t even hear the colleague’s objection, and so does not possess the counterevidence that it furnishes.) Goldman (1992) seems to concur with Kornblith’s assessment that the physicist’s belief that p is defeated in this case. But it’s not clear that ARP can deliver this verdict, since the reliable process involved in listening to his colleague’s testimony wouldn’t be entirely internal—it would, in an important sense, involve engaging in further research.

I won’t press this objection, however, since I don’t find Kornblith’s case entirely convincing. While I certainly concur that there’s something epistemically defective about the physicist’s overall state, it’s not clear to me that this defect defeats his belief that p . (It seems plausible to me that a person can have justified beliefs, even though those beliefs are sustained, at least in part, by a variety of epistemic vices.)

What's more, Harry has a very high opinion of Unger's 1975 time-slice: whenever it occurs to Harry that Unger would advise him (Harry) to suspend judgement about p , this causes Harry to become chagrined and suspend judgement about p . So if Harry had used his Unger Predictor, he would have come to believe SUSPEND: \langle Unger would advise me (Harry) to suspend judgement regarding TREE \rangle . This would, in turn, have caused Harry to suspend judgement regarding TREE.

Sufficiency predicts that Harry's belief in TREE is defeated at t . After all, his Unger Predictor is a reliable process: it systematically produces true beliefs about the advice of Unger's 1975 time-slice. It's also an *internal, cognitive* process that's available to him at t . So there's a reliable, internal belief-forming process that's available to Harry at t , which, if it had been used by Harry, would have resulted in him not believing TREE at t .

However, this verdict about the case strikes me as wrong: my intuition is that Harry's belief in TREE is justified at t . After all, at t , Harry isn't thinking about Unger, or entertaining any sceptical doubts; indeed, we can stipulate he hasn't engaged in sceptical ruminations in a very long time. At t , it seems that he has every reason to think that there's a tree in front of him and no good reason to suspend judgement. The mere fact that *if* he were to engage in reflection about what Unger would think about his situation, he would suspend judgement regarding TREE shouldn't deprive his belief of justification.

Thus I take it that Thinking About Unger is a counterexample to Sufficiency, and hence ARP. Is there any way of repairing ARP to get around this counterexample? When faced with this case, it's natural to feel that there's something fishy about the connection between Harry's Unger Predictor and his belief in TREE. And thus it's natural to suspect that if we can flesh out this worry—if we can articulate in what sense this connection is fishy—we'll be able to revise ARP in a way that avoids generating the prediction that Harry's belief in TREE is defeated.

So what exactly is fishy about the connection between Harry's Unger Predictor and his belief in TREE? It seems to me that there are two natural hypotheses. The first is that *causal indirectness* is to blame. If Harry were to use his Unger Predictor, this wouldn't directly cause him to suspend judgement about TREE. The only immediate effect of his

Unger Predictor would be a belief in SUSPEND; the actual suspension of judgement regarding TREE would occur as a further, downstream effect.

A second, closely related hypothesis is that a subject matter mismatch is the culprit. According to this hypothesis, Harry's Unger Predictor doesn't produce beliefs about the right sort of subject matter: it produces beliefs about what Unger would advise, not about whether there are trees in front of Harry.

These two hypotheses suggest two different ways of trying to repair ARP. The first hypothesis suggests a version of ARP that includes a causal directness requirement:

ARP*: S's belief that p is defeated at t iff there's some reliable or conditionally reliable process available to S at t which, if it had been used by S in addition to the process actually used, would have directly resulted in S's not believing p at t .

The second hypothesis suggests a version of ARP that includes a subject matter requirement. A natural way of spelling out this requirement is to insist that the alternative process must produce a doxastic attitude towards p :

ARP:** S's belief that p is defeated at t iff there's some reliable or conditionally reliable process X available to S at t such that:

1. The output of X is a doxastic attitude towards p
2. If X had been used by S in addition to the process actually used, it would have resulted in S's not believing p at t .

Will either ARP* or ARP** work? I think not. To see why, imagine that Harry had used his Unger Predictor, and hence had come to suspend judgement regarding TREE. Now, presumably there is *some* process that directly results in this suspension of judgement; it's just a two-component process. The first component is his Unger Predictor, which outputs a belief in SUSPEND; the second component is a process we can call his 'Unger Implementer': it takes as input a belief in SUSPEND, and produces suspension of judgement regarding TREE. Call this two-component process Harry's 'Unger Emulator' (depicted in fig. 1.1).

Now, Harry's Unger Emulator is a reliable process. After all, it produces no beliefs; *a*

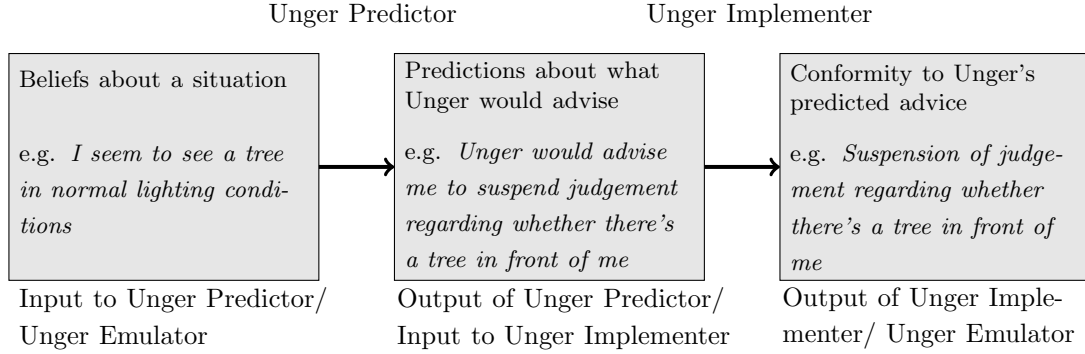


Figure 1.1: Unger Emulator

fortiori, it doesn't produce any false beliefs. So it seems Harry has an alternative reliable belief-forming process available to him (his Unger Emulator) which, if he had used it in addition to the process actually used (visual perception), would have *directly* resulted in him suspending judgement about TREE. Hence, ARP* delivers the result that his belief in TREE is defeated at t .¹⁴ What's more, the output of his Unger Emulator is a doxastic attitude towards TREE (suspension of judgement is, after all, a doxastic attitude); hence ARP** also delivers the verdict that his belief in TREE is defeated at t .

One might object to my claim that Harry's Unger Emulator is a reliable belief-forming process. There are at least two ways one might try to motivate this objection. First, one could point out that since the Unger Emulator doesn't produce any beliefs, it doesn't count as a *belief-forming* process at all. Second, one could dispute the assumption that failure to produce false beliefs is sufficient for reliability. On a sufficiently nuanced conception of reliability, suspending judgement on a truth will count as worse than believing a truth.¹⁵ Since the Unger Emulator suspends judgements on all truths, it's not a particularly reliable process.

¹⁴Could one insist that, strictly speaking, the process that would have *directly* resulted in his suspension of judgement wouldn't have been his Unger Emulator; rather, it would have been just the second stage of the Unger Emulator (the Unger Implementer)? Even if this view could be motivated (which I doubt), it will be of no help to the proponent of ARP*; after all, the Unger Implementer produces no beliefs either, hence it too is reliable.

¹⁵The best candidates for nuanced conceptions of reliability along these lines come from the literature on scoring rules, where the discussion is usually couched in terms of credences rather than the tripartite distinction between belief, suspension, and disbelief. See, for example, Joyce 1998; Gibbard 2007; Moss 2011. (Goldman also discusses such scoring rules in places (Goldman and Shaked 1991, Goldman 1999, Goldman 2010), though he presents such scoring rules as measures of "degrees of truth possession" or "veritistic value" rather than reliability.)

But even if we grant this objection, we can amend the case. Imagine an epistemologist—call her Shmunger—whose scepticism is much more circumscribed than that of Unger’s 1975 time-slice. Shmunger has lots of true beliefs about math, physics, history—you name it. Shmunger is only a sceptic about trees: in any given situation, she suspends judgement on whether trees are present.

Suppose, furthermore, that Shmunger wants others to share her doxastic attitudes: she wants others to hold her (entirely accurate) beliefs about math, physics, etc.; she also wants everyone to share her arboreal scepticism. Now we can reproduce the counterexample using Shmunger:

Thinking About Shmunger: Harry’s twin Larry has an highly reliable Shmunger Predictor: given any scenario, Larry’s Shmunger Predictor reliably predicts what Shmunger would advise Larry to believe in that scenario. Larry’s Shmunger predictor is also part of a Shmunger Emulator: whenever Larry predicts that Shmunger would advise Larry to take up a particular doxastic attitude towards a proposition, he invariably takes up that attitude. Now Larry sees a tree in normal lighting conditions. Since Larry isn’t thinking about Shmunger, he comes to believe TREE: ⟨There is a tree in front of me⟩. But if he had used his Shmunger Emulator, he wouldn’t have done so.

Clearly, Larry’s Shmunger Emulator is a belief-forming process: it produces lots of beliefs on sundry topics. What’s more, the beliefs in question are always true. Consequently, it seems to qualify as a reliable belief-forming process, even given a suitably nuanced conception of reliability. If Larry had used this process, it would have *directly* resulted in Larry’s suspending judgement about TREE. So ARP* predicts that Larry’s belief that there’s a tree in front of him is defeated. And since the output of Larry’s Shmunger Emulator is a doxastic attitude towards TREE, ARP** delivers the same prediction.

So it seems that the most natural ways of revising ARP to get around the counterexamples to Sufficiency prove unsuccessful.¹⁶ And it’s worth noting that nothing hinges on the

¹⁶For those who suspected that some sort of subject matter mismatch was to blame, a natural next attempt at patching ARP would be the following:

ARP*:** S’s belief that p is defeated at t iff there’s some process X available to S at t such that:
(i) X is reliable about p -related matters

details of Thinking About Unger/Shmunger; we can construct many more counterexamples with the same general structure. Here's the formula:

Step 1: Describe a case where an agent S has a *prima facie* justified belief that *p*.

Step 2: Endow S with a dormant reliable belief-forming process X that meets the following conditions:

1. If S were to use X, the output of X would be suspension of judgement about *p*
2. The mere availability of X does not, intuitively, give S a good reason to cease believing *p*.

What's more, even if we can revise ARP in a way that avoids every case of this form, ARP won't be in the clear: as we're about to see, there are also counterexamples to Necessity. Any satisfactory revision of ARP will have to deal with these as well.

1.5 A Counterexample to Necessity

Here's a counterexample that shows that ARP fails to articulate necessary conditions for defeat:

Job Opening: Masha tells Clarence that her department will have a job opening in the fall. Clarence believes Masha; assuming that Masha is usually reliable, Clarence's belief

(ii) If X had been used by S in addition to the process actually used, it would have resulted in S's not believing *p* at *t*.

If we coupled ARP*** with a nuanced conception of reliability, we could capture the intuition that Larry's belief in TREE is justified. After all, though his Schmunger Emulator is generally reliable, it isn't reliable about TREE-related matters: it systematically recommends suspending judgements about the presence of trees, hence isn't reliable about their presence.

However, proponents of this account face the difficult task of specifying what counts as a "*p*-related matter." (As an anonymous referee at *The Philosophical Quarterly* helpfully observed, this could be viewed as a "second round" of the Generality Problem.) If in Thinking About Shmunger we take a "TREE-related matter" to include any proposition about the presence of trees (as the foregoing paragraph suggests), then the envisioned fix won't be able to handle cases of even more circumscribed scepticism. (Imagine Shmunger isn't a wholesale sceptic about the presence of trees; indeed, she's generally quite reliable about the presence of trees. She's only a sceptic about the presence of, say, birch trees when viewed from such-and-such an angle.) To deal with such cases, we'd need to give a more restrictive account of "TREE-related matters." But the more we restrict what counts as a "*p*-related matter", the more difficulty we'll have capturing intuitions about mundane cases of defeat. For instance, in Consuela and the Vase, we'd like to say that Consuela's belief in RED is defeated even though her interlocutor isn't reliable about whether the particular vase Consuela is looking at is red.

counts as *prima facie* justified. Sometime later, Clarence speaks with the head of Masha's department, Victor, who informs him that the job search was cancelled due to budget constraints. Now suppose that Clarence harbors a deep-seated hatred of Victor that causes him to disbelieve everything that Victor says; what's more, no amount of rational reflection would rid Clarence of this inveterate distrust. Consequently, he continues to believe that there will be a job opening in the fall.

Intuitively, Clarence's belief that there will be a job opening in the fall is defeated.¹⁷ But it doesn't seem that Necessity delivers this result. Consider: is there some alternative reliable process available to Clarence such that, if Clarence had used it in addition to the process he actually used (namely, trusting Masha's testimony), he wouldn't have believed that there will be a job opening in the fall? It doesn't seem so. After all, we've stipulated that Clarence has an intractable distrust of Victor—one that no amount of reflection or therapy could dislodge. So it doesn't seem that there's any process available to Clarence that takes Victor's testimony as input and outputs a fairly high credence in that testimony.

One might insist that Clarence has a very general process (GP) available to him, which takes any testimony *t* he's received as input and produces an increased credence in *t* as output. It's just that Clarence isn't capable of employing this process for every value of *t*. But even if we grant that he has such a process available to him, this won't be enough to rescue Necessity. In order to rescue Necessity, the following counterfactual would need to be true:

If Clarence had used GP, he wouldn't have believed there will be a job opening in the fall.

But given the way the case is set up, this counterfactual is false. After all, Clarence has only received two pieces of testimony regarding the potential job opening: Masha's and Victor's. He can't plug Victor's testimony into GP, and plugging Masha's testimony into GP was what caused him to believe that there will be a job opening in the first place.

¹⁷It would be different if Clarence had good reason to deem Victor's testimony unreliable, but in our scenario this isn't the case.

Of course, one could part ways with Goldman and opt for a less restrictive conception of the belief-forming processes ‘available’ to an agent. Someone who takes this approach could point out that Clarence is capable of asking other people whether the department will have a job opening in the fall; if Clarence had pursued such inquiries, he would have presumably acquired independent evidence that the job search was canceled. Our defender of Necessity could then insist that this suffices for Clarence to have an alternative reliable process available to him which, if used in addition to the process actually used, would have resulted in him abandoning the belief that there will be a job opening in the fall.

However, there are at least two problems with this maneuver. First, we can set up the case in such a way that Clarence’s independent investigations would have been fruitless: simply stipulate that, despite persistent inquiries, Clarence would have been unable to unearth any further evidence regarding the potential job opening. Second, a process reliabilist who makes this maneuver will open herself up to a host of additional counterexamples to Sufficiency. For instance, she’ll fall prey to Kvanvig’s counterexample (discussed in §1.4): if engaging in further inquiry counts as a belief-forming process available to an agent, Sufficiency entails that S’s belief that p will be defeated whenever there’s some reliable cognizer with whom S is acquainted who believes $\neg p$.

Here too nothing hinges on the details of case; it’s easy to cook up similar examples. Here’s the recipe:

Step 1: Describe a case where an agent S has a *prima facie* justified belief that p .

Step 2: Stipulate that S receives strong evidence e that p is false.

Step 3: Stipulate that, due to hatred, prejudice, or just some psychological quirk, S is cognitively incapable of responding appropriately to e , and hence persists in believing p .

Every case we cook via this recipe will put defenders of Necessity in a similar bind. For every such case, defenders of Necessity may be tempted to insist that even though S is unable to respond appropriately to e , S is capable of pursuing further inquiries into the truth of p ; had she done so, she would have unearthed some further evidence e^* that would have caused her to suspend judgement on p . This, they may be tempted to insist, is all that it takes for S to satisfy Necessity, given a suitably liberal conception of ‘availabil-

ity.’ But broadening our conception of availability will invariably lead from frying pan to fire: any such broadened conception of availability will give rise to legions of Kvanvig-style counterexamples to Sufficiency.

1.6 Looking Forward

Where does this leave us? We’ve seen that process reliabilists need an account of defeat (§2). What’s more, the standard process reliabilist account of defeat—ARP—falls victim to counterexamples. Thus process reliabilists must either supply an alternative reductive account of defeat or renounce their reductive ambitions altogether.

While I’ve been focused on process reliabilism, I should note that the difficulties I’ve raised aren’t *just* a problem for process reliabilism. They are, at the very least, problems for any attempt to give a reductive ‘Two Step’ theory of justification (where a Two Step theory of justification is any theory that analyzes justification in terms of the satisfaction of two conditions: a *prima facie* justification condition and a ‘No Defeaters’ condition). For example, consider the dogmatist view that whenever it perceptually appears to S that *p*, S is *prima facie* justified in believing *p* (Pryor 2000). Or consider the more general ‘phenomenal conservative’ position that whenever it seems to S that *p*, S is *prima facie* justified in believing *p* (Huemer 2007). Clearly, the mere fact that it perceptually appears to S that *p* is not sufficient for S to be *ultima facie* justified in believing *p*. To see this, we need look no further than the case of Consuela: it perceptually appears to her that there’s a red vase in front of her, but—as we’ve seen—she isn’t *ultima facie* justified in believing this. Thus, any such view needs a ‘No Defeaters’ clause as well. And if such a view is intended to be reductive, then it will need to explain defeat in non-epistemic terms.

At this point, some may argue that what we’ve been missing is any appeal to the notion of *evidence*. As Comesaña (2010) and Goldman (2011) observe, the notion of *evidence* is conspicuously lacking from traditional reliabilist accounts. Perhaps this notion gives us just what we need to arrive at a counterexample-free view of *ultima facie* justification.

In the next two sections, I consider this idea in depth. I argue that in order for this move to do the desired work, we’ll need to cash out the notion of *evidence possession* in

non-epistemic terms. Moreover, I argue that this is no easy task. Indeed, it is arguably just as difficult as explicating the notions of *justification* and *defeat* in non-epistemic terms.

1.7 Appealing to Evidence

There are two main strategies for enlisting the notion of *evidence* into a theory of justification. The first is to use evidence to explain defeat.¹⁸ For example, one could propose an account of defeat along the following lines:

Evidential Account of Defeat: S's belief that p is defeated at t iff S's total evidence at t doesn't support p at t .

One could then combine the Evidential Account of Defeat with some account of *prima facie* justification to deliver a Two Step theory of *ultima facie* justification. For example, we could combine the Evidential Account of Defeat with a process reliabilist account of *prima facie* justification, resulting in a hybrid of process reliabilism and evidentialism. (For recent defenses of a synthesis of process reliabilism and evidentialism, see Comesaña 2010; Goldman 2011; and Tang forthcoming.)

A second strategy is to abandon the 'Two Step' structure altogether. Rather than trying to explain epistemic justification in terms of a combination of (i) an account of *prima facie* justification, (ii) a 'No Defeaters' condition, perhaps we should instead give a 'One Step' analysis of epistemic justification in terms of evidence. This would result in a 'pure' evidentialist view. While there are different ways of spelling out such a view, here is a characteristic statement:

Evidentialism about Justification: S is (*ultima facie*) justified in believing p at t iff S's total evidence supports p at t .¹⁹

(Note that evidentialism, thus formulated, is a thesis about propositional justification. Typically, evidentialists go on to analyze doxastic justification in terms of propositional

¹⁸See e.g., Conee and Feldman 2005 for the view that defeat is best characterized in terms of evidence.

¹⁹See Conee and Feldman 2004, 2008. (For formulations of evidentialism as a supervenience thesis, see Fantl and McGrath 2002; Weatherson 2005.)

justification plus the basing relation—e.g., S’s belief that p is doxastically justified iff (i) it’s based on body of evidence e that propositionally justifies believing p , (ii) there’s no more inclusive body of evidence e' possessed by S that doesn’t propositionally justify S in believing p (Conee and Feldman 2004). Whether such an analysis of doxastic justification will work is controversial: see Turri 2010 for putative counterexamples; see Goldman 2011 for the view that any such account will need to be supplemented with constraints on the processes that produce the target belief. For present purposes, I will set such worries aside.)

There is an obvious hurdle to using either of these strategies in service of a reductive account of justification: the notion of *possessing evidence* is clearly an epistemic notion. (Arguably, the notion of *supporting a proposition* is an epistemic notion as well.) Thus, evidentialists with reductive ambitions need to provide a reductive account of *evidence possession*.²⁰ But providing such an account is easier said than done. To see why, let us start by considering a plausible constraint on a theory of evidence possession:

Evidence-Justification Link: If S has p as evidence at t , then S is (*ultima facie*) justified in believing p at t .

To see this plausibility of the Evidence Possession-Justification Link, suppose it were false. Then we should expect it to be perfectly fine to say things like:

- (1) ? The detective’s evidence entails that the butler did it, but the detective isn’t justified in believing the butler did it.
- (2) ? Given John’s evidence, there’s no possibility that Jane is in the office. Still, John shouldn’t think that Jane isn’t in the office.

²⁰Goldman (2011) articulates a closely related worry: namely, that evidentialists owe us an account of *evidence*. As Goldman observes, the threat of circularity looms. After all, some characterize evidence as whatever justifies belief (Kim 1988: 290-291). And another popular approach—discussed in detail in Chp.3—identifies one’s evidence with one knows (Williamson 2000a). If one conjoins this approach with the traditional view that knowledge can only be analyzed in terms of justification, then evidentialists will be saddled with a circular account of justification.

My challenge here is subtly different: evidentialists with reductive ambitions owe us not just an account of evidence; they also owe us an account of evidence *possession*. Even if the former can be explicated in non-epistemic terms, it’s not clear that the latter can be. To see this, consider Goldman’s own suggestion that evidence is best understood in terms of reliable indication: e is evidence for p iff e reliably indicates p (Goldman 2011). Clearly, something can be evidence in this sense, even if no one ‘possesses’ this evidence. (Consider, e.g., the number of rings on an undiscovered tree, which reliably indicate the tree’s age.) Thus even if Goldman’s suggestion gives an adequate reductive account of unpossessed evidence, reductive evidentialists will still need to tell us what is required for an agent to ‘possess’ a body of evidence.

But such discourses strike me as incoherent, or at least borderline incoherent. This suggests that our ordinary intuitions about evidence possession conform to the Evidence-Justification Link.²¹

Our argument for the Evidence-Justification Link need not rest entirely on intuitions about evidence possession. We can instead appeal to the connection between evidence possession and evidential support. Consider the following plausible principle:

Evidence-Support Link: If S has p as evidence at t , then S's total evidence supports p at t .

Most leading accounts of evidential support validate this principle. Consider, for example, probabilistic theories of support, according to which the degree to which S's evidence supports p at t is given by a probability function Pr that reflects the probability of p conditional on S's total evidence e at t (e.g., Williamson 2000a). On any such theory, if $p \in e$, then $\text{Pr}(p|e) = 1$, and so S's evidence will support p to the greatest degree possible. Or consider modal accounts of evidential support, according to which S's evidence e supports p iff in every normal world where e obtains, p obtains (e.g., Smith 2010). Clearly if $p \in e$, then in every world where e obtains, p obtains. And so modal accounts also validate the Evidence-Support Link.

But if evidence is to explain defeat and/or *ultima facie* justification, one can't accept the Evidence-Support Link while denying the Evidence-Justification Link. After all, suppose the Evidence-Justification Link were false. In particular, suppose there is some subject S who has some proposition q as evidence, even though S's belief in q is defeated. Given the Evidence-Support Link, S's evidence supports q . Given either the Evidential Account of Defeat or Evidentialism about Justification, it follows that S belief in q isn't defeated—which contradicts our initial supposition. Thus those who seek to explain justification in terms of evidence must either accept the Evidence-Justification Link or they must develop a non-standard account of evidential support that invalidates Evidence-Support Link.

Giving a reductive account of evidence possession that validates the Evidence-Justification Link is no trivial task. In the following section, I consider three initially promising strate-

²¹For further arguments for the Evidence-Justification Link, see Beddor 2015a.

gies for developing a reductive account of evidence possession and argue that each fails to undewrite the Evidence-Justification Link. It remains to be seen whether some alternative reductive strategy fares better.

1.8 Three Attempts to Reduce Evidence Possession

The first reductive strategy is *mentalist*: it seeks to characterize evidence possession in terms of an agent’s non-factive mental states.²² How exactly do an agent’s non-factive mental states determine the facts about evidence possession? Mentalists are often cagey when it comes to answering this question. But here are some natural accounts that can be gleaned from the literature:

Conscious Beliefs: S has p as evidence iff S consciously believes p .²³

Experiences and Memories: S has p as evidence iff p is the content of one of S’s perceptual experiences or apparent memories.²⁴

Module Outputs: S has p as evidence iff p is the output of one S’s modules.²⁵

One could also combine these proposals in various ways. For example, Schroeder (2011b) defends the view that S has p as evidence iff S has a “presentational attitude” towards p , where a “presentational attitude” is any attitude that represents its content as true. Schroeder’s discussion makes it clear that he takes both belief and perceptual experience to be presentational attitudes.

None of these accounts are consistent with the Evidence-Justification Link. To see this, consider again Consuela’s belief in RED. Even after receiving her interlocutor’s testimony, she consciously believes RED; moreover, on many views of perceptual content, RED is also the content of her visual experience. And, given the assumption that vision is a module,

²²See e.g., Lewis 1996; Conee and Feldman 2004. (The ‘non-factive’ qualification is meant to exclude non-reductive views such as Williamson 2000a, which identifies evidence with knowledge, but takes knowledge to be a mental state.)

²³See Feldman 1988 for the view that S has p as evidence at t iff S is thinking of p at t .

²⁴For a classic account of evidence in terms of perceptual experiences and apparent memories, see Lewis 1996.

²⁵See Weatherson 2009. (Weatherson understands modules in roughly the sense of Fodor 1983.)

it is also the output of one her modules. Nonetheless, she isn't *ultima facie* justified in believing RED.²⁶

A second reductive strategy attempts to analyze evidence possession in terms of beliefs that track the truth in certain ways—e.g., beliefs that are safe or sensitive.²⁷ Such views also run afoul of the Evidence-Justification Link. To see this, consider the following simple truth-tracking approach:

Safe Belief: S has p as evidence iff S safely believes p

where S safely believes $p =_{def}$ in all nearby worlds where S believes p on the same basis, p is true.

Note that Consuela's belief in RED could well be safe: we can stipulate that in all sufficiently nearby worlds where she uses vision to arrive at the belief that there is a red vase in front of her, this belief is true. Still, her belief in RED isn't *ultima facie* justified.²⁸

A third strategy is to try to use the resources of virtue epistemology to give a reductive account of evidence. For example, one could hold that S has p as evidence iff S is disposed

²⁶Given this, some mentalists may be inclined to retreat to a more austere conception of evidence. For example, rather than taking an agent's evidence to be the contents of her perceptual experiences, some may be inclined to hold that her evidence consists entirely in propositions about how things appear to her at a time. On this proposal, Consuela's evidence doesn't include RED; rather, it includes a proposition of the form: \langle It appears to me that RED \rangle . Perhaps, mentalists may suggest, any such proposition is guaranteed to be *ultima facie* justified.

However, this maneuver faces at least two difficulties. First, it's doubtful that beliefs about appearances are immune to defeat. Suppose the vase's color is a borderline shade of red: it is red, but a slight change in its hue would make it orange. And suppose a trustworthy oracle informs Consuela that she regularly makes mistakes about how color properties appear to her in borderline cases, and that in this very case, she's making a mistake. Arguably, her belief that it appears to her that there's a red vase in front of her is defeated.

Second, this route seems to lead towards skepticism. If our evidence only consists in an extremely austere class of propositions—propositions about appearances—it's unclear how we could be justified in forming beliefs about the external world. For further development of this idea that an austere conception of evidence leads to skepticism, see Williamson 2000a.

²⁷Here one could either give a 'direct' reduction and simply define S's evidence as S's beliefs that possess some truth-tracking property T , or one could give an 'indirect' reduction and analyze evidence in terms of some further epistemic state—e.g., knowledge—and then analyze this further epistemic state in terms of beliefs that possess T .

²⁸One might try to avoid this result by insisting on a particularly fine-grained way of typing Consuela's basis for believing RED. However, this is tantamount to reverting back to the Typing Defense, which, as we've already seen (§1.2), faces difficulties. For further criticism of this maneuver, see Lasonen-Aarnio 2010. (Note that while Aarnio's conclusion is largely congenial to mine, she seems to think that defeat poses a special problem for externalists. If my arguments in this chapter are correct, this is mistaken: defeat is a general problem for all reductive accounts of justification, be they externalist or internalist.)

to competently assent to p .²⁹ Those who prefer a factive conception of evidence could hold that S has p as evidence iff S aptly believes p , where S 's belief that p is apt iff it is successful (i.e., true) because of the exercise of a belief-forming competence.³⁰

In order to evaluate this strategy, we need to say more about epistemic competences. One option is to understand epistemic competences in epistemically loaded terms, for example, as dispositions to form justified beliefs, or as dispositions to know (Millar 2009; Miracchi 2015). But clearly this is tantamount to abandoning a reductive epistemology. Alternatively, we could opt for a reductive conception of epistemic competences, identifying them with dispositions to form true beliefs (where perhaps these dispositions are relativized to particularly methods of belief formation). But if we go this route, then once again we're forced to deny the Evidence-Justification Link. After all, Consuela's belief in RED is the result of vision operating in good lighting conditions, which presumably qualifies as a disposition to form true beliefs, and hence as an epistemic competence (at least on this reductive understanding of epistemic competences). Moreover, the truth of her belief in RED is the result of an exercise of this epistemic competence. And so it seems that on the reductive virtue epistemological approach, she has this proposition as evidence, even though she is not *ultima facie* justified in believing it.³¹

1.9 Conclusion

In this chapter, I have argued that leading attempts to provide a reductive account of justification—as well as their natural successors—flounder when it comes to dealing with

²⁹See Sylvan and Sosa (2014) for an account of epistemic reasons along these lines. (Note that Sylvan and Sosa are noncommittal on whether all epistemic reasons can be identified with evidence.)

³⁰See Sosa 2007. There are a variety of slightly different ways of cashing out this idea. For example, we might follow Turri (2011b) and Sosa (2010, 2015) and instead focus on whether the belief is true because of the “manifestation” of an epistemic competence. The differences between these formulations, while interesting in their own right, should not matter for my purposes.

³¹Sosa (2010, 2015) distinguishes three features of a competence such as archery: (i) the constitution or seat (which is retained even when the archer is asleep), (ii) an appropriate inner condition or “shape” (e.g., being awake and sober), (iii) an appropriate situation or environment (e.g., good lighting conditions). Whether this “triple S” account of competences qualifies as reductive depends on how we understand the notion of “appropriate” inner conditions and situations in the case of epistemic competences. Once again we face a dilemma. On the one hand, we could understand these in an epistemically loaded fashion—e.g., by defining “appropriate” conditions and situations as those in which the competence will produce justified beliefs—in which case the resulting account will fail to be reductive. Or we can try to cash them out in non-epistemic terms, in which case it's not clear how we'll be able to deliver the right results in cases like Consuela and the Vase.

cases of epistemic defeat. Such attempts either fall victim to counterexample or invoke some unanalyzed epistemic notion (e.g., *defeat*, *evidence*) and hence fail to be reductive.

Of course, I have no proof that a reductive account of justification is impossible. Thus one option at this point is to press forward, and try to give some better reductive account of justification. However, I will confess that it is far from clear to me what such an account will look like.

A second option is to abandon our aspirations of giving a reductive account of justification, but to insist that the problems that arose for reductive accounts of justification do not undermine reductive accounts of other epistemic notions—e.g., knowledge and evidence. However, I think there is reason to find this response both unappealing and unconvincing. To see why this response is unappealing, recall our original motivation for giving a reductive epistemology. We started with the observation that *all* epistemic facts seem to supervene on non-epistemic facts. This in turn motivated Reducibility: the thesis that all epistemic facts are reducible to the non-epistemic facts (§1.1). If we can only reduce *some* epistemic notions to the non-epistemic, this won't suffice to underwrite the Reducibility.

Moreover, the second option is unconvincing in light of the arguments in §§1.7-1.8. There I argued that possessing *p* as evidence entails being *ultima facie* justified in believing *p* (§1.7). Furthermore, I argued that this entailment imposes a significant constraint on any adequate account of evidence, and indeed that many leading accounts in the literature fail to satisfy it (§1.8). Similar remarks apply to knowledge. According to a venerable—though by no means uncontested—view, knowing *p* entails being *ultima facie* justified in believing *p*. If this is right, then this clearly places an important constraint on any adequate analysis of knowledge. Thus the problems for reductive accounts of justification threaten to generalize to reductive accounts of other epistemic notions.

A third—and, in my eyes, more attractive—option is to conclude we've been pursuing the wrong *sort* of reductive analysis. To put the idea in very general terms, we should continue to try to analyze justification and other epistemic notions in non-epistemic terms, but we should pursue analyses that take a rather different form than those pursued so far. In the next chapter, I explore one way of developing this option in more detail.

Chapter 2

Towards a Reductive Metaepistemology

2.1 Introduction

The analyses of justification that we explored in Chapter One were couched in the material mode. That is, they were of the form:

- (3) S is justified in believing p iff...

In this chapter, I explore what happens if we take the path of semantic ascent, and couch reductive analyses in the formal mode. That is, I will focus on epistemic linguistic expressions and try to analyze such expressions in non-epistemic terms. Thus I will be pursuing analyses of the form:

- (4) “S is justified in believing p ” is true (relative to parameters $\sigma_1 \dots \sigma_n$) iff...

where the right-hand side of the biconditional is couched in non-epistemic terms.¹

¹Of course, it is often harmless to semantically ascend and descend. For example, it would be harmless to recast Simple Process Reliabilism in the formal mode: “S’s belief that p is justified” is true at t iff S’s belief that p at t is the result of a reliable belief-forming process. However, some expressions block any straightforward attempt to semantically descend. Context-sensitive expressions are a familiar example. Consider a simple semantic analysis of “here” talk: “ x is here” is true in a context of utterance c iff x is located in the vicinity of the speaker in c . There is however, no straightforward way of restating this analysis in the material mode. According to the account developed in this chapter, epistemic expressions are similar to “here” talk in this regard.

How should we understand the difference between analyses that take the form of (3) and those that take the form of (4)? In the ethical domain, philosophers often distinguish between first-order ethics and metaethics. First-order ethics attempts to specify the conditions under which an act (or person, or state of affairs) instantiates some moral property. For example, utilitarianism is a first-order theory: it says that an action is right iff it maximizes utility. By contrast, metaethics investigates the nature of our moral thought and talk (as well as the nature of the properties that such thought and talk picks out). At least one important branch of metaethics—what we might call, ‘semantic metaethics’—focuses on moral language, and attempts to specify the meanings of moral linguistic expressions (e.g., terms such *right*, *wrong*, *good*). By analogy, moving from (3) to (4) can be viewed as moving from first-order epistemology to metaepistemology.²

What would a reductive metaepistemology look like? The particular reductive strategy I will explore in this chapter takes its cue from the metaethics literature. In metaethics, there is a rich tradition of analyzing moral expressions in terms of agents’ desire-like attitudes.³ One attractive feature of this approach is its reductive potential: given the assumption that the desire-like attitudes can be understood without recourse to normative notions, this approach delivers an analysis of moral expressions in non-normative terms. While this ‘attitudinal’ approach has long been popular in metaethics, relatively few have tried to extend this approach to metaepistemology.⁴ And of the few who have, relatively few have tried to work out the semantic details in a convincing fashion. In this chapter, I explore how this can be done: I analyze a variety of epistemic linguistic expressions in terms of agents’ desire-like attitudes.

Here is the plan for the chapter. In §2.2 I develop a preliminary account of a distinctive desire-like attitude that is involved in epistemic evaluations—an attitude that I call ‘epistemic approval.’ In §2.3 I show how we can enlist this attitude into semantic service. Specifically, I develop an attitudinal semantics for a variety of epistemic expressions, focus-

²My claim is not that the only task of metaepistemology is to give analyses of epistemic linguistic expressions; I only claim that this is a task for metaepistemology—what we might call, ‘semantic metaepistemology’.

³See e.g., Ayer 1936; Stevenson 1944, 1963; Blackburn 1984, 1998; Brogaard 2008; Dreier 1990; Gibbard 1990, 2003, 2008; Harman 1975; Horgan and Timmons 2006; Ridge 2006; Prinz 2007; Egan 2012; Starr forthcoming, among many others.

⁴Though the recent surge of interest in epistemic expressivism marks an important exception: see e.g. Gibbard 2003; Ridge 2007; Field 2009; Kappel 2010; Chrisman 2007, 2012a; Grajner 2015.

ing on epistemic evaluations involving deontic modals (e.g., “You ought to believe in global warming”) and epistemic justification ascriptions. (I also briefly discuss how the approach can be extended to knowledge ascriptions and evidence possession ascriptions.) Along the way, I suggest how parallel attitudinal semantics can be developed for a variety of moral expressions—including moral evaluations involving deontic modals and moral justification ascriptions. In §2.4 I advertise the advantages of the resulting metaepistemology. The main advantage, from the present perspective, is its reductive potential: epistemic expressions can be understood in terms of the psychological attitude of epistemic approval, which can be understood in non-epistemic terms. Moreover, the reduction in question nicely parallels attitudinal reductions in metaethics, thereby offering a unified strategy for reducing the normative to the non-normative. But there are further advantages to boot. In particular, this approach explains some of the striking commonalities between epistemic and ethical expressions—for example, the fact that both are *value-laden* and the fact that both are connected to motivation. §2.5 concludes by discussing the relations between metaepistemological reductions and first-order epistemological reductions.

2.2 Epistemic Approval

Folk psychology contains a variety of closely related conative attitudes: e.g., desire, preference, intention, and approval. And each of these conative attitudes can in turn be divided into various subspecies. Take, for instance, approval. We can distinguish between approving of something *all things considered* and approving of something from a particular point of view.

An example may help. Imagine an avaricious business mogul who reads Singer’s “Famine, Affluence, and Morality.” Try as he might, he cannot identify any flaws in Singer’s argument. And so he becomes reluctantly convinced that ought to give donate most of his fortune to charity. As a result, he has some pro-attitude towards a state of affairs in which he donates most of his income to charity: when he contemplates this state of affairs, it strikes him as being morally commendable. We can say that he ‘morally approves’ of this state of affairs: he approves of it from a moral point of view.

Still, he may not approve of it all things considered. Suppose our mogul has a soft spot for the finer things in life, and would be loath to forego the various luxuries to which he has grown accustomed. From a purely prudential point of view, he does not approve of giving away most of his fortune. What's more, we may well imagine that for our mogul, the prudential point of view trumps the moral: all things considered, he most approves of hanging onto his riches.

In addition to approving of something from a moral point of view, it seems we sometimes approve of something from an epistemic point of view. Suppose that Shelly thinks that, given the evidence available to Fred, he ought to believe in global warming. It seems to follow that she has some pro-attitude towards a state of affairs in which Fred believes in global warming. As we can put it, she 'epistemically approves' of this state of affairs.

Epistemic approval can come apart from both moral approval and all things considered approval. To see this, consider mediocre Mel, who believes he is a talented musician, despite all the evidence to the contrary. Suppose, moreover, that Mel's belief in his musical abilities is central to his wellbeing: it is a source of pride and hope; without it, he would despair. Shelly might well epistemically disapprove of Mel's belief that he is a talented musician: she might think that, given all his evidence, it is irrational for him to hold this belief. Nonetheless, she does not morally disapprove of this belief: she does not think there is anything immoral with his believing as he does. And she might not all things considered disapprove of this belief. Indeed, she might think that, all things considered, it is best for Mel to go on believing as he does.

So far I have introduced the notions of *moral approval* and *epistemic approval* through examples. But what is it, exactly, to approve of something from the moral or the epistemic point of view?⁵ One might well worry we cannot answer this question without relying on moral or epistemic notions. If this turns out to be the case, then even if we can analyze moral and epistemic linguistic expressions in terms of moral and epistemic approval, the

⁵In the metaethics literature, the problem of identifying some distinctively moral conative attitude is known as the 'moral attitude problem.' (The label is due to Miller (2003), who argues against a number of putative solutions. For relevant discussion, see Kauppinen 2010; Köhler 2013; Björnsson and McPherson 2014.) While this is often raised as a problem for noncognitivism, it seems a version of the moral attitude problem—and its epistemic analogue, which we might call the 'epistemic attitude problem'—arises for anyone who thinks that distinctive conative attitudes are typically involved in moral and epistemic judgment.

reductive *bona fides* of the resulting analysis will be in question.

However, I think this pessimism is unwarranted. There are at least three promising strategies for developing reductive accounts of moral and epistemic approval. I will mention each in turn.

The first strategy is to analyze different subspecies of approval as dispositions towards all things considered approval—dispositions that have distinctive bases. In the case of moral approval, the basis might be a *pro tanto* desire for fairness or the wellbeing of sentient creatures. According to this proposal, our avaricious mogul has a disposition to all things considered approve of donating his fortune to charity—a disposition that’s based in a *pro tanto* desire to alleviate suffering. However, this disposition is masked by prudential considerations—considerations that prevent him from all things considered approving of donating his fortune.

In the case of epistemic approval, the distinctive basis might be a *pro tanto* desire for the attainment of truth and the avoidance of error. According to this proposal, Shelly is disposed to all things considered approve of Mel abandoning his belief in his musical abilities—a disposition that’s based in a *pro tanto* desire for error avoidance. However, this disposition is masked by her concern for Mel’s wellbeing.

A second strategy is to analyze different subspecies of approval in terms of their causal connections with both (i) other psychological attitudes, (ii) behaviors.⁶ For example, one might hold that moral approval is distinguished by the fact that is intimately connected with the ‘reactive attitudes’ (Strawson 1974) such as guilt, shame, and blame. For example, if S morally approves of ϕ -ing in circumstances c , then it seems natural to expect S to be disposed to feel some degree of guilt if she recognizes that she is in c , but nonetheless doesn’t

⁶One way of embedding this strategy in a more general framework for understanding psychological states is to adopt the picture in Lewis 1972, 1994, according to which folk psychology can be regarded as a theory, consisting in various platitudes about the causal connections between different mental states. To define a particular mental state term t (e.g., belief, desire), we can then collect up the platitudes and ‘Ramsify’ over them—that is, conjoin them in a long sentence ϕ , replace every occurrence of t in ϕ with a variable x , and then define t as the unique entity that makes the resulting sentence true.

Of course, applying this strategy to moral and epistemic approval may be somewhat dicey, since ‘moral approval’ and ‘epistemic approval’ aren’t familiar terms of folk psychology. However, I’m hopeful that my initial description of moral and epistemic approval struck a chord with readers: the states are familiar, even if folk psychology doesn’t contain commonly accepted names for them. If this is right, then the Ramsification strategy for moral and epistemic approval may indeed hold promise. For further discussion of how the Ramsification strategy might help solve the moral attitude problem, see Köhler 2013.

attempt to ϕ . And it would likewise be natural to expect S to be disposed to blame others who are in c , but do not attempt to ϕ .⁷

In a similar vein, we might propose that epistemic approval is distinguished by the fact that it is intimately connected with patterns of belief formation and revision. To see the appeal of this proposal, imagine a character—call him ‘Ted’—who epistemically disapproves of atheism. According to Ted, various considerations (the ontological argument and the like) make it epistemically obligatory to be a theist. We’d naturally expect Ted to be a theist: were we to learn that he’s a staunch atheist, his epistemic disapproval of atheism notwithstanding, we would be surprised. And if at some later time Ted changes his mind and comes to epistemically disapprove of theism (due to the problem of evil and the like), we would presumably revise our expectations concerning Ted’s theism (or lack thereof): we would now expect Ted to be an atheist or an agnostic. More generally, it seems that if S epistemically approves of having some doxastic attitude D towards p in circumstances c , we expect S to be disposed to adopt D towards p upon coming to learn that she is in c . According to the proposal under consideration, this intimate connection with our doxastic attitudes is partially constitutive of epistemic approval.⁸

Yet another strategy is to analyze moral and epistemic approval in terms of the functions they serve. Borrowing from Gibbard (1990)’s account of ‘accepting a norm’, one might suggest that the function of moral approval is to foster coordination among our attitudes, expectations, and behaviors. Similarly, one might follow Dogramaci (2012) and hold that the function of epistemic approval is to foster coordination in our belief-forming methods (coordination which, Dogramaci argues, helps ensure that testimony is reliable).⁹

⁷Kauppinen (2010) also invokes the reactive attitudes in his account of moral approval/disapproval, but he puts them to rather different use. According to Kauppinen, moral disapproval is “an attitude of disapproval that characteristically results from impartially placing oneself in the shoes of those primarily affected by the action and sharing their negative reactive attitudes.” Kauppinen’s proposal thus uses the reactive attitudes in specifying a characteristic basis for moral approval/disapproval. For discussion of Kauppinen’s proposal, see Björnsson and McPherson 2014.

⁸Some authors deny that epistemic akrasia—that is, cases where one does not believe in accordance with how one thinks one epistemically ought to believe—is possible. (See e.g., Hurley 1989; Pettit and Smith 1996.) This strong position strikes me as implausible: consider, for example, the undergraduate who, after reading Descartes’ *Meditations*, becomes convinced that it’s irrational to believe in the existence of an external world, but nonetheless goes around believing in tables and chairs. (For other examples, see Greco 2014.) But while cases of epistemic akrasia seem to be possible, they also seem highly unusual: our default expectation is that people will be epistemically enkratic. According to the present proposal, this observation reveals something important about the nature of epistemic approval.

⁹Dogramaci (2012) focuses on the function of our epistemic evaluations, which for him are various lin-

I've briefly sketched three strategies for analyzing subspecies of approval: (i) analyzing them as dispositions towards all things considered approval—dispositions which are individuated by their bases, (ii) analyzing them in terms of their causal connections with various states and behaviors, (iii) analyzing them in terms of the functions they serve. Each of these strategies strikes me as promising. And there is no need to choose between them: perhaps all three are required for a full picture of moral and epistemic approval. For example, here's a toy example of the sort of definitions of moral and epistemic approval that might result from combining them:

Toy Theory of Moral Approval:

Moral approval =_{def} the psychological relation R that meets the following conditions:

1. Anyone who stands in R towards p is disposed to all things considered approve of p —a disposition that's based in a *pro tanto* desire for fairness and/or wellbeing.
2. *Ceteris paribus*, a psychologically normal agent who stands in R towards p will be disposed to feel some degree of remorse upon realizing that she could have easily brought p about, but failed to do so.
3. The function of R is to coordinate various expectations, attitudes, and behaviors.

Toy Theory of Epistemic Approval:

Epistemic approval =_{def} the psychological relation R that meets the following conditions:

1. Anyone who stands in R towards p is disposed to all things considered approve of p —a disposition that's based in a *pro tanto* desire for the attainment of truth and avoidance of error.
2. *Ceteris paribus*, a psychologically normal agent who stands in R towards p will be disposed to bring her doxastic policies in accordance with p .
3. The function of R is to coordinate various belief-forming methods.

guistic acts. However, it would be natural to extend his account to the underlying conative attitudes that—according to the present proposal—we give voice to when we perform these linguistic acts.

These toy definitions are almost certainly too simple. However, they should at least give the reader a sense for the sort of definitions of moral and epistemic approval that might emerge from a combination of the three strategies I’ve outlined. And they should license optimism about the possibility of characterizing moral and epistemic approval without reliance on any unanalyzed moral or epistemic notions. If such a characterization can be given, then epistemic and moral approval can be used as the foundation for a reductive metaepistemology and metaethics.

2.3 An Attitudinal Semantics

Equipped with an account of epistemic approval, I now explore how we can put this attitude to work in our semantics for various epistemic linguistic expressions.

I start by looking at epistemic evaluations involving deontic modals (e.g., *ought*, *should*, *must*, *may*). Consider, for example:

- (5) a. You ought to believe in global warming.
- b. You shouldn’t take a stand on whether the butler did it—you should wait until we gather more evidence.

Both (5a) and (5b) have natural readings on which they make epistemic evaluations (as opposed to say, moral or prudential evaluations). The reading in question can be made explicit using phrases such as *in view of the evidence* and the somewhat more cumbersome, *from an epistemic point of view*:

- (6) a. In view of the evidence, you ought to believe in global warming.
- b. From an epistemic point of view, you shouldn’t take a stand on whether the butler did it.¹⁰

I start by focusing on epistemic evaluations involving deontic modals because deontic modals have been well-studied in the linguistics literature, and there is by now a fairly

¹⁰These epistemic evaluations involving deontic modals should not be conflated with what are sometimes called, ‘epistemic modals’—that is, uses of *might*, *may*, *ought* and *must* to convey whether some proposition is compatible with (or entailed by) some agents’ epistemic state. (For example, “It might rain later”.) I offer an analysis of epistemic modals in Chp.3.

standard (though not entirely uncontroversial) semantic analysis of such expressions (due to Kratzer 1981, 1991, 2012). This makes them a natural starting place: I can piggyback on the standard semantic analysis of deontic modals in developing my attitudinal approach. After laying out my treatment of such expressions (§§2.3.1), I turn to consider how this approach can be extended to other epistemic linguistic expressions, such as justification ascriptions, evidence ascriptions, and knowledge ascriptions (§§2.3.2–2.3.3).

2.3.1 An Attitudinal Semantics for Deontic Modals

Deontic *Ought* (First Pass)

According to the standard treatment, modals are sentential operators.¹¹ Thus (5a) is analyzed as:

- (7) Ought (You believe in global warming).

The next step in the standard analysis is to relativize the truth of modals to two contextually determined parameters: a modal base and an ordering source. The modal base is a function f from a world to a set of accessible worlds.¹² Exactly which worlds f delivers will vary with context: in some cases, the modal base will deliver the set of all worlds (W); in some cases, it will deliver the set of worlds that resemble the actual world in certain relevant respects; in some cases, it will deliver the set of worlds compatible with the speaker's epistemic state, etc. The ordering source is a function g from a world to a set of propositions that serve to 'rank' the worlds in the modal base along some dimension: it says which worlds in the modal base are 'better' (along the relevant dimension) than others. Exactly what sort of 'betterness' is in play will also vary with context: some ordering sources will care about which worlds are *morally* better than others; some will care about which worlds are *more normal* than others; some will care about which worlds conform to certain laws or regulations better than others, etc.

¹¹The idea that modals function as sentential operators is widely endorsed in the literature, though see Schroeder 2011a for the view that at least some modals have a reading that relates agents directly to actions. For a reply to Schroeder, see Chrisman 2012b.

¹²Officially, Kratzer takes the modal base to deliver a set of propositions; we get a set of accessible worlds by intersecting these propositions.

Putting all of this together, modals quantify over the ‘best’ worlds in the modal base. For example, *ought* is analyzed as a universal quantifier: “Ought ϕ ” says that all of the best worlds in the modal base are worlds where the prejacent (ϕ) is true.

One standard way of formalizing this is to assume that $g(w)$ induces a preorder (a reflexive and transitive relation) $\succeq_{g(w)}$ over W . This preorder corresponds to the “at least as good as” relation (determined by $g(w)$): for any worlds v and u , $v \succeq_{g(w)} u$ iff v is at least as good as u according to $g(w)$. We can use this preorder to define the ‘best’ worlds in the modal base as follows:

$$\text{BEST}_{g(w),f(w)} = \{v : v \in f(w) \text{ and } \forall u \in f(w): v \succeq_{g(w)} u\}$$

We can then formulate Kratzer’s semantics for *ought* as follows:

Kratzer’s Semantics for *Ought*:

$$\llbracket \text{Ought } \phi \rrbracket^{c,f,g,w} = 1 \text{ iff } \forall w' \in \text{BEST}_{g(w),f(w)} : \llbracket \phi \rrbracket^{c,f,g,w'} = 1$$

To illustrate, Kratzer’s semantics predicts that (5a)/(7) will be true in a context of utterance iff all of the best worlds in the modal base are worlds where the addressee believes in global warming. To capture the reading where (5a)/(7) makes a distinctly epistemic evaluation, we will presumably use an ordering source that ranks the worlds in the modal base from an epistemic point of view. On this reading, (5a)/(7) says that all of the epistemically best worlds in the modal base are worlds where the addressee believes in global warming.

Kratzer’s semantics is an impressive apparatus. To mention one particularly attractive feature, it offers a unified semantics for modals in natural language: it explains why the same modal expression (e.g., *ought*) can be used to make very different claims—for example, an epistemic evaluation, a moral evaluation, or a prudential/bouletic evaluation. Each of these different claims will be modeled using a different ordering source, which ranks the worlds in the modal base along a different dimension—the epistemic point of view, the moral point of view, the prudential point of view, etc.

Despite this, it’s important to note that Kratzer never gives us a substantive, philosophically illuminating account of the ordering sources. In particular, she never answers the following question:

Q: What it is for a particular set of worlds to be epistemically best?

This is all well and good for Kratzer’s purposes: after all, she wants a story about modals that’s general enough to handle all flavors of modals—deontic, bouletic, teleological, etc. But those who are interested in a reductive metaepistemology will want something more: they’ll want a substantive answer to Q. Here’s where the attitudinal framework comes in: it says that *certain agents’ states of epistemic approval determine what count as the epistemically best worlds*.

Here’s a simple way of implementing this idea. Let e_S^w denote a subject S’s *epistemic approval profile* at a world w : a complete specification of all the states of affairs that S epistemically approves of at w , together with various degrees to which S epistemically approves of these states of affairs. We can use an agent’s epistemic approval profile to induce an ordering $\succeq_{e_S^w}$ over W : for any worlds v and u , $v \succeq_{e_S^w} u$ if and only if v is at least as good, from the point of view of S’s epistemic approval profile at w , as u . We now let the ordering source used for the epistemic evaluation reading of deontic *ought* be the epistemic approval profile of some contextually selected subject (e.g., the speaker). According to this proposal, (5a)/(7) will be true (on the reading where it makes an epistemic evaluation) if and only if all of the worlds in the modal base that are ranked highest by the contextually selected subject’s epistemic approval profile are worlds where the addressee believes in global warming.

This strategy can be easily extended to handle moral evaluations using deontic modals, such as:

(8) You (morally) ought to save the drowning child.

All that’s required is that we posit a *moral approval profile* (m_S^w)—a specification of the states of affairs that S morally approves of at w (together, perhaps, with the degree to which S morally approves of those states of affairs). Then we can suggest that when it comes to moral evaluations involving deontic modals, the ordering source is some contextually selected subject’s moral approval profile. According to this proposal, (8) will be true if and only if all of the worlds that are ranked highest by the contextually selected subject’s moral approval profile are worlds where the addressee saves the drowning child. Thus the current

approach offers to give us a reductive semantics for both epistemic and moral evaluations involving deontic *ought*.

Potential Revisions and Refinements

There are number of ways in which one could modify my reductive semantics for epistemic and moral evaluations involving deontic *ought* without departing from the guiding idea. Here I'll mention a few such modifications.

On the simple approach sketched above, the context of utterance always selects a particular subject whose epistemic approval profile determines the ordering when it comes to epistemic evaluations involving deontic *ought*. But this could be rejected: perhaps the context sometimes selects a group of agents, whose aggregate state of epistemic approval determines the ordering when it comes to epistemic evaluations involving deontic *ought*. (Perhaps these are the agents are the speakers' interlocutors; perhaps they are the agents who share the speaker's values; perhaps they are the agents that the speaker aspires to resemble in certain respects, etc.) And there is no need to assume that there will always be a unique group selected by the context of utterance. Borrowing a move from von Fintel and Gillies (2011), we could say that sometimes an utterance of an epistemic evaluation involving a deontic modal doesn't assert a single determinate proposition; instead, it "puts in play" a cloud of propositions. Thus an utterance of (5a)/(7) might put in play (i) the proposition that all of the worlds in the modal base ranked highest by the speaker's state of epistemic approval are worlds where the addressee believes in global warming, (ii) the proposition that all of the worlds in the modal base ranked highest by the conversational participants are worlds where the addressee believes in global warming, etc.

Another potential modification is to inject a relativist element into our semantics. Thus far, our semantics has been contextualist: we have only relativized the semantic values of deontic modals to contexts of utterance (and parameters determined by the context of utterance). According to relativists about deontic modals, further relativization is needed. Relativists insist that even if we hold fixed the context of utterance, a sentence containing a deontic modal could have different truth-values relative to different assessors (that is,

different individuals who are assessing the sentence for truth or falsity).¹³ For instance, consider a sentence such as:

- (9) Everyone ought to believe in the external world.

According to relativists, (9) might be true relative to an assessor who epistemically approves of everyone believing in the external world (say, G.E. Moore), but false relative to a skeptic, who epistemically disapproves of so believing.

A simple way of implementing this relativist idea in the present framework is to take our indices to be not just worlds, but *centered worlds*: that is, world, assessor pairs. When it comes to epistemic evaluations, the ordering source is not simply a function from a world to the epistemic approval profile of some subject; rather, it's a function from a w, a pair to the epistemic approval profile of some subject (or some group of subjects). In the simplest case, the subject in question will just be the assessor. This gives us the following semantics:

Relativist Semantics for *Ought*:

$$\llbracket \text{Ought } \phi \rrbracket^{c,f,g,w,a} = 1 \text{ iff } \forall w' \in \text{BEST}_{g(w,a),f(w)} : \llbracket \phi \rrbracket^{c,f,g,w',a} = 1$$

According to this semantics, (9) will be true relative to some index w, a iff all of the worlds in the modal base that are ranked highest by the epistemic approval profile of the w, a -determined subject are worlds where everyone believes in the external world. Those attracted to moral relativism could easily extend the present treatment to moral uses of deontic *ought*: the moral ordering source will be a function from a world, assessor pair to the moral approval profile of some subject(s).¹⁴

What would decide between our original contextualist semantics and this relativist variant? Thus far most of the dispute between contextualists and relativists has centered on disagreement data. Consider, for instance, a dispute such as:

- (10) a. *Moore*: Everyone ought to believe in the external world.

¹³For relativist treatments of various expressions, see Lasersohn 2005; Egan et al. 2005; Egan 2007, 2010, 2012, 2014; Stephenson 2007a,b,c; MacFarlane 2011, 2014.

¹⁴Note that this semantics for deontic *ought* is only relativist about the ordering source, not the modal base. However, it's certainly open to take the modal base to vary with the assessor as well. For relevant discussion, see Kolodny and MacFarlane 2010.

- b. *Skeptic*: $\left\{ \begin{array}{l} \text{That's false!} \\ \text{Nuh uh!} \end{array} \right\}$ We ought to suspend judgment.

Relativists will typically allege that a simple version of contextualism—say, one that holds that the ordering source is always the epistemic approval profile of the speaker—has difficulty accounting for the intuition that the skeptic disagrees with Moore’s claim (since presumably the skeptic doesn’t disagree with the claim that Moore epistemically approves of believing in the external world). By contrast, relativists predict that Moore’s utterance was false for the skeptic. Hence relativists claim that they have an easier time capturing the disagreement data.

Whether relativists have a genuine advantage over contextualists when it comes to accounting for disagreement is highly controversial, and will not be settled here.¹⁵ My point is simply that those who are persuaded by the arguments for relativism can easily implement my semantics in a relativist framework.

One further potential modification is worth mentioning. In the metaethics literature, the most popular version of an attitudinal semantics is neither contextualism nor relativism, but rather expressivism. Expressivism about some sentence ϕ holds that utterances of ϕ do not serve to “represent” anyone (a speaker, or group of agents, etc.) as having certain conative attitudes; rather, they directly express these attitudes. Applied to epistemic evaluations involving deontic modals, the view would be that the epistemic evaluation reading of a sentence such as (9)/(10a) directly expresses epistemic approval of everyone believing in the external world.

Stated thus, expressivism is rather vague, and there are various ways of cashing it out more precisely. One option, advocated by Yalcin (2012a), construes expressivism as a pragmatic thesis. Applied to epistemic evaluations involving deontic modals, the view is that the conventional conversational effect of asserting a sentence such as (9)/(10a) is to try to get one’s interlocutors to epistemically approve of a state of affairs in which everyone believes in the external world. For the sake of perspicuity, I will adopt this conception of expressivism in what follows.

¹⁵For contextualist rejoinders to the argument from disagreement, see de Sa 2008; Cappelen and Hawthorne 2009; Sundell 2011; Pearson 2013; Plunkett and Sundell 2013, among others.

Even if expressivism is a pragmatic thesis, it places important constraints on this semantics. One way to semantically implement this thesis is to take semantic contents to be sets of world, epistemic approval profile pairs.¹⁶ Importantly, our epistemic approval states here will not be indexed to any particular agents; rather, they are specifications of states of affairs that are epistemically approved of to various degrees—specifications that could model the epistemic approval state of multiple agents, provided those agents epistemically approve of the same states of affairs to the same degrees. These epistemic approval profiles can serve as ordering sources, determining what count as the ‘best’ worlds in the modal base. Letting e stand for any such epistemic approval profile, we can state the semantics as follows:

Expressivist Semantics for *Ought* (First Pass):

$$\llbracket \text{Ought } \phi \rrbracket^{c,f,w,e} = 1 \text{ iff } \forall w' \in \text{BEST}_{e,f(w)} : \llbracket \phi \rrbracket^{c,f,w',e} = 1,$$

$$\text{where } \text{BEST}_{e,f(w)} = \{v : v \in f(w) \text{ and } \forall u \in f(w) : v \succeq_e u\}$$

According to this approach, the content of an assertion of e.g., (9)/(10a) will be the set of w, e pairs such that e assigns highest ranking to those worlds in the modal base where everyone believes in the external world. And since the goal of an assertion is to get one’s interlocutors to accept its content, the goal of asserting (10a) is to get one’s interlocutors to epistemically approve of believing in the external world.

Those attracted to moral expressivism could adopt a natural extension of the present approach. The obvious move is to take indices to be sets of ordered triples of worlds, epistemic approval profiles, and moral approval profiles. The ordering for a deontic modal can be provided by either the epistemic approval profile in the index or by the moral approval profile. When the deontic modal is used to make an epistemic evaluation, the ordering will be provided by the epistemic approval profile (as before). When the modal is used to make a moral evaluation, the ordering will be provided by the moral approval profile. (Whether a particular use of deontic modal is making a moral or epistemic evaluation will be left to context.) Let m stand for a moral approval profile, the semantics can be stated as follows:

¹⁶This strategy is similar to Gibbard’s suggestion that the contents of normative sentences are sets of world, norm pairs (1990) (or world, hyperplan pairs (2003)).

Extended Expressivist Semantics for *Ought*:

$$\llbracket \text{Ought } \phi \rrbracket^{c,f,w,e,m} = 1 \text{ iff } \forall w' \in \text{BEST}_{c,e,m,f(w)} : \llbracket \phi \rrbracket^{c,f,w',e,m} = 1.^{17}$$

The viability of expressivism is fiercely debated in the current literature, and it is not my goal to settle this issue here. Rather, my aim is simply to advertise the flexibility of my attitudinal metaepistemology: it can be implemented in a contextualist, relativist, or expressivist setting. This flexibility strikes me as a mark in its favor.

Other Modals

So far I have sketched an attitudinal semantics for epistemic and moral evaluations involving deontic *ought*. But this approach can easily be extended to other modals.

Ought and *should* are typically classified as weak necessity modals, in contrast to strong necessity modals (*must*, *have to*, *need to*).¹⁸ As their name suggests, weak necessity modals are weaker than strong necessity modals: (11) entails (12), but not *vice versa*.

$$(11) \quad \left\{ \begin{array}{c} \text{Must} \\ \text{Has to} \end{array} \right\} \phi.$$

$$(12) \quad \left\{ \begin{array}{c} \text{Should} \\ \text{Ought to} \end{array} \right\} \phi.$$

As evidence for this, consider:

$$(13) \quad \text{Tony} \left\{ \begin{array}{c} \checkmark \text{ should} \\ \# \text{ has to} \end{array} \right\} \text{mow the lawn, but he doesn't have to.}$$

(13) is felicitous when the first conjunct contains a weak necessity modal, but contradictory when this is replaced by a strong necessity modal.

For further evidence, compare:

¹⁷This is just one way of developing an expressivist semantics. Another promising strategy is to implement expressivism using the resources of dynamic semantics. Charlow (2015) and Starr (forthcoming) take precisely this line. For example, Starr takes deontic modals to perform tests on conversational states, where a conversational state c is an ordered pair of a set of worlds c_s and a desirability ordering \succ over c_s . A state c passes the test imposed by “Ought ϕ ” iff the most desirable worlds (according to \succ) in c_s are ϕ -worlds. Starr’s proposal can easily be developed in a way that implements my attitudinal metaepistemology. All we’d need to do is replace the desirability ordering \succ with an epistemic approval ordering (provided by the conversational participants’ epistemic approval profile).

¹⁸See Sloman 1970; Horn 1972, 1989; Harman 1993; McNamara 1996; von Fintel and Iatridou 2005, 2008; Copley 2006; Portner 2009; Lassiter 2011; Chrisman 2012c, among others.

(14) ?? Tony ought to mow the lawn; in fact, he should mow the lawn.

(15) ✓ Tony ought to mow the lawn; in fact, he $\left\{ \begin{array}{l} \text{must} \\ \text{has to} \end{array} \right\}$ mow the lawn.

(14) is infelicitous, unlike (15). A natural explanation of the infelicity of (14) is that it's infelicitous to reinforce a sentence ϕ with a sentence that is obviously entailed by ϕ .¹⁹ If *must* were equivalent to *should*, (15) would be similarly infelicitous.²⁰

How should we model the meanings of strong necessity modals? This is controversial, but here's a natural picture, inspired by Sloman (1970) and developed recently by von Fintel and Iatridou (2005, 2008). Suppose that in addition to distinguishing between the very best worlds and the suboptimal worlds, the ordering source distinguishes between the *acceptable* worlds and the *unacceptable* worlds in the modal base. Intuitively, a world is acceptable as long as it is *good enough* from the point of view of the ordering source, where what's good enough may—in at least some contexts—fall short of the best. We then say that strong necessity modals universally quantify over the acceptable worlds in the modal base:

$$\llbracket \text{Must } \phi \rrbracket^{c,f,g,w} = 1 \text{ iff } \forall w' \in \text{ACCEPTABLE}_{g(w),f(w)} : \llbracket \phi \rrbracket^{c,f,g,w'} = 1$$

Clearly, the best worlds will always be a subset of the acceptable worlds, but not *vice versa*: there will often be acceptable worlds that still fall short of being optimal. And so this view correctly predicts that (11) entails (12), but not *vice versa*.

Suppose we accept this way of modeling the distinction between weak and strong necessity modals. Then this could easily be implemented using my attitudinal metaepistemological framework. The idea would be that certain agents' states of epistemic approval distinguish not only between the best worlds and the suboptimal worlds; they also distinguish between the acceptable worlds and the unacceptable worlds. (One way of developing this idea would be to say that a world is acceptable as long as it's consistent with the things

¹⁹See Sadock 1978; Stanley 2008; Littlejohn 2011.

²⁰There is an interesting question as to whether the distinction between weak and strong necessity modals is cross-linguistically robust. At least some other languages, such as German, also lexicalize the distinction between weak and strong necessity modals (e.g., *sollen* (weak) vs. *müssen* (strong)). von Fintel and Iatridou (2008) argue that a number of languages express weak necessity by augmenting a strong necessity modal with counterfactual morphology (e.g., French *Il devrait faire la vaisselle* (weak) vs. *Il doit faire la vaisselle* (strong)).

that the relevant agents approve of to a sufficiently high degree.) This gives us the resources to capture epistemic evaluations that are expressed using strong necessity modals, e.g.:

- (16) You must not believe contradictions.

On the simple contextualist version of our semantics, (16) will be true if and only if there are no acceptable worlds (from the point of view of the contextually selected agent's epistemic approval profile) in the modal base in which the addressee believes contradictions.²¹

This approach can also be extended to handle epistemic evaluations involving expressions of permission, e.g.:

- (17) Given the evidence, it's (epistemically) permissible for the jurors to believe that the defendant is guilty.

It's commonly thought that expressions of permission are the duals of strong necessity modals:

Permission–Strong Necessity Duality:

Permitted ϕ iff $\neg(\text{Must } \neg\phi)$

As evidence for Permission–Strong Necessity Duality, note that the following sounds contradictory:

- (18) # You must not believe contradictions. But you may believe contradictions.

If we conjoin Permission–Strong Necessity Duality with our analysis of strong necessity modals, we get the view that expressions of permission are existential quantifiers over the acceptable worlds in the modal base:

²¹It's somewhat difficult to find any clear cases where what you epistemically must believe comes apart from what you epistemically ought to believe. (This is closely related to the question of whether there is epistemic supererogation—see e.g., Tidman 1996; Hedberg 2014.) But even if there are no such cases, it seems we should recognize at least a conceptual or semantic difference between strong and weak necessity modals, and hence between the epistemic evaluations that use the former and those that use the latter. (Compare: on some versions of utilitarianism, you morally ought to ϕ in exactly those circumstances where you morally must ϕ —namely, the circumstances in which ϕ -ing maximizes utility. But even if some such version of utilitarianism is true, there is still a difference in felicity between “You ought to give all of your disposable income to charity, but you don't have to” and “You must give all of your disposable income to charity, but you don't have to.” We'd like our semantics for modals to explain this difference.)

$$\llbracket \text{Permitted } \phi \rrbracket^{c,f,g,w} = 1 \text{ iff } \exists w' \in \text{ACCEPTABLE}_{g(w),f(w)} : \llbracket \phi \rrbracket^{c,f,g,w'} = 1$$

Given our attitudinal interpretation of ordering sources, this delivers the result that (16) is true iff there is at least one acceptable world in the modal base (from the point of view of the contextually selected agent's epistemic approval profile) in which the jurors believe that the defendant is guilty.

2.3.2 An Attitudinal Semantics for Justification Ascriptions

Thus far I have been focusing on epistemic evaluations involving deontic modals. In this section, I consider two strategies for extending my attitudinal metaepistemology to justification ascriptions such as:

- (19) The jurors are justified in believing the defendant is guilty.
- (20) Consuela is justified in believing there's a red vase in front of her.

The first strategy analyzes justification ascriptions in terms of deontic modals, thereby enabling us to leverage our analysis of deontic modals into an analysis of justification ascriptions. The second strategy analyzes *justified* as a degree expression, and analyzes justification ascription in terms of a sufficiently high degree of epistemic approval.

The Deontological Approach

According to what are sometimes called *deontological approaches to justification*, we can analyze *justified* in deontic terms.²² How, exactly, does the analysis go? Proponents of deontological approaches are not always clear on this point.²³ However, three natural options present themselves, corresponding to the three types of deontic modals that we have identified thus far. The options are: (i) the Weak Necessity View, which takes justification ascriptions to be equivalent to epistemic evaluations involving weak necessity modals,

²²By now there's a large literature on deontological approaches to justification. See e.g. Alston 1988; Plantinga 1993: chp.1; the papers in Steup 2001; Littlejohn 2012: chp.1.

²³Thus far, most of the literature on deontological approaches to justification has focused on the question of whether *any* form of a deontological approach to justification is in principle viable. In particular, most of the literature focuses on Alston's (1988) objection, according to which deontological approaches to justification entail an implausible form of doxastic voluntarism. For present purposes, I'll set Alston's challenge aside. For relevant discussion, see Kim 1994; Chuard and Southwood 2009; Nottelman 2013.

(ii) the Strong Necessity View, which takes justification ascriptions to be equivalent to epistemic evaluations involving strong necessity modals, (iii) the Permissive View, which takes justification ascriptions to be equivalent to epistemic evaluations involving permissive expressions. To illustrate, these three views will analyze (20) as (21a)-(21c), respectively:

- (21) a. Consuela (epistemically) ought to believe there's a red vase in front of her.
 b. Consuela (epistemically) must believe there's a red vase in front of her.
 c. Consuela is (epistemically) permitted to believe there's a red vase in front of her.

In addition to these options, Beddor (forthcoming) argues that there is a fourth possibility that has been overlooked by the literature—namely, that justification ascriptions are to weak necessity modals as expressions of permissions are to strong necessity modals. To elaborate this possibility, let us posit a *faultlessness* operator, defined as the dual of weak necessity modals:

Faultlessness–Weak Necessity Duality:

Faultless ϕ iff $\neg(\text{Ought } \neg\phi)$

When conjoined with our semantics for deontic *ought*, Faultlessness–Weak Necessity Duality leads to the view that expressions of faultlessness are existential quantifiers over the best worlds in the modal base:

$$\llbracket \text{Faultless } \phi \rrbracket^{c,f,g,w} = 1 \text{ iff } \exists w' \in \text{BEST}_{g(w),f(w)} : \llbracket \phi \rrbracket^{c,f,g,w'} = 1$$

According to what we can call the ‘Faultlessness View’, *justified* functions as a faultlessness operator. And so on this approach, (20) is equivalent to:

- (22) $\neg(\text{Consuela (epistemically) ought not believe that there's a red vase in front of her}).$

In other words: at least one of the epistemically best worlds in the modal base is a world where Consuela believes there's a red vase in front of her.²⁴ For a visual comparison of these four deontic statuses, see fig. 2.1.

²⁴My notion of faultlessness shouldn't be conflated with more familiar hypological notions, such as *blamelessness*. (Here I borrow the term ‘hypological’ from Zimmerman (2002), who uses it refer to notions relating

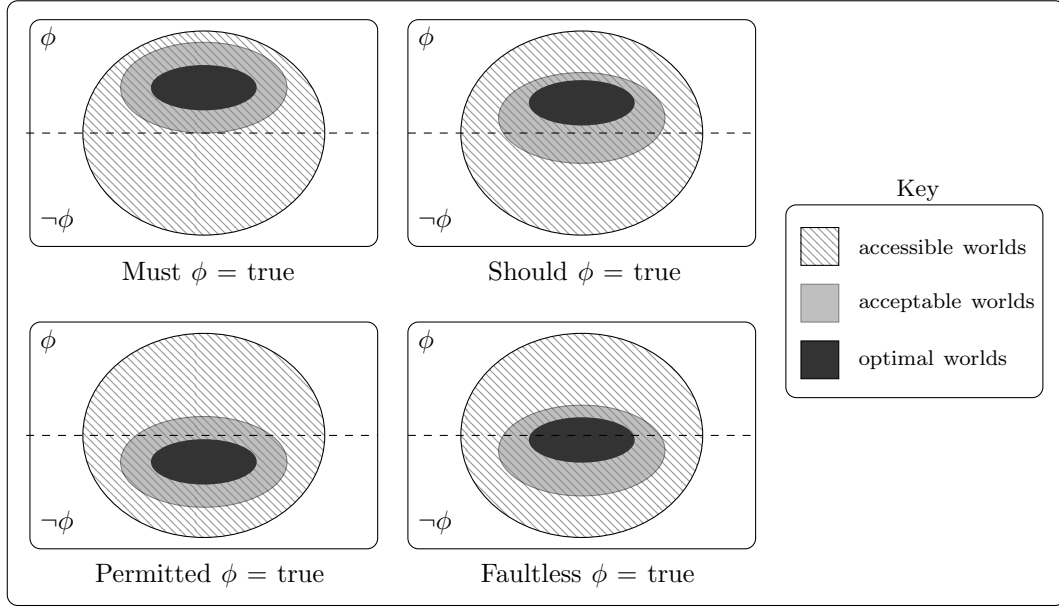


Figure 2.1: Four Deontic Operators.

Which of these four versions of a deontological approach is the most promising? This is a difficult question, and not one that I will try to fully settle here. However, it is worth raising a couple of pieces of data that may help decide among our four candidates. (I develop these arguments in more detail in Beddor forthcoming.) The first concerns the entailments of *justified* under negation. The following inference seems valid:

- (23) a. S isn't justified in believing p . \Rightarrow
 b. S (epistemically) shouldn't believe p .

For instance, if I tell you that Consuela isn't justified in believing there's a red vase in front of her, then I seem to have committed myself to the claim that she shouldn't believe that there is a red vase in front of her.

This suggests that *justified* does not function as a necessity modal (either weak or strong). After all, the corresponding inference involving a necessity modal is invalid:

to responsibility.) Suppose the speed limit is 45 mph. Suppose that Shelly is driving 50 mph, but that she has a reasonable false belief that she's doing 40 (perhaps because her speedometer has unforeseeably malfunctioned). In this case, she's doing something that she ought not do (relative to the normative standard provided by the speed limit), hence her action is not faultless in our sense. But presumably her action is blameless. More generally, a distinction between faultlessness and blamelessness will arise whenever it's possible to have reasonable false beliefs about whether one is doing what one should do.

- (24) a. $\neg \left\{ \begin{array}{c} \text{Ought} \\ \text{Must} \end{array} \right\} \phi. \not\Rightarrow$
- b. $\left\{ \begin{array}{c} \text{Ought} \\ \text{Must} \end{array} \right\} \neg \phi.$

For example, it's not the case that I should eat Cornflakes for breakfast; eating Raisin Bran would be just as good. But it doesn't follow that I shouldn't eat Cornflakes for breakfast.

This leaves us with either the Permissive View or the Faultlessness View, both of which validate the inference from (23a) to (23b).²⁵ Is there any data that favors one over the other? One argument in favor of the Faultlessness View is that instances of the following schema sound very odd:

- (25) S is justified in believing p , but S (epistemically) should suspend judgment on p .

To my ears—and those of many that I've polled—the following sound infelicitous:

- (26) ?? Kwame is justified in believing it will rain, but Kwame should suspend judgment on whether it will rain.
- (27) ?? Kendra is justified in believing the restaurant is open, but Kendra should suspend judgment about whether the restaurant is open.²⁶

The Faultlessness View explains this sense of incoherence. According to the Faultlessness View, the first conjunct of (25) says that it's not the case that S shouldn't believe p , which contradicts the second conjunct. By contrast, the Permissive View has a harder time explaining this data. As we've seen, weak necessity modals don't entail strong necessity

²⁵ According to the Permissive View, (23a) is analyzed as: ' \neg (S is permitted to believe p)', which simplifies to: 'S must not believe p ', which in turn entails (23b). According to the Faultlessness View, (23a) is analyzed as: ' \neg (\neg (S should not believe p))', which simplifies to (23b). (Note that while both the Permissive View and the Faultlessness View validate this entailment, the Faultlessness View is the logically strongest semantics for *justified* that does so. After all, on the Faultlessness View, (23a) is *equivalent* to (23b). Since negation reverses logical strength, any stronger semantics for *justified* would entail that (23a) is weaker than (23b).)

²⁶ A caveat: these sentences are only infelicitous if the modal in the second conjunct (*should*) is interpreted as making an epistemic (rather than, say, a bouletic or prudential) evaluation. For example, if Kwame will gain a million dollars if he suspends judgment on whether it will rain, we can access a coherent reading of (26) on which the second conjunct says that, given his interest in gaining a million dollars, he should suspend judgment. For my purposes, the important data point is that these sentences are infelicitous when the modal is interpreted as making a distinctly epistemic evaluation. For further discussion of this issue, see Beddor forthcoming.

modals, and permission is the dual of strong necessity modals. So ‘Permitted ϕ ’ is compatible with ‘Ought $\neg\phi$ ’—both will be true in a situation where ϕ obtains in at least one acceptable world, but no optimal world. For a concrete example, consider the following case due to Elizabeth Harman:

Feedback Quandary: Amanda is a philosophy professor who has a two-year-old daughter. It is 11pm. Amanda receives an email from her undergraduate student Joe, with a third draft of a paper that’s due tomorrow at noon. She has already commented on the first two drafts. Joe is struggling in the class, but she can tell he is on the verge of some kind of breakthrough. If Joe fails the class, he will lose his scholarship and have to drop out of school. It would take half an hour to read the draft and write the comments, and Amanda is tired. Her daughter will wake up early. Amanda realizes that she is not morally obligated to spend the thirty minutes to give Joe comments, but nevertheless she deliberates about whether to do it. Upon reflection, Amanda thinks, ‘I should do it!’ (Harman 2016)

Harman offers the following verdict about this case:

(28) Amanda should give Joe comments, but it would be permissible not to.²⁷

This judgment strikes me as plausible. At the very least, it is perfectly *coherent*. And so proponents of the Permissive View are faced with a puzzle as to why instances of (25) sound worse than (28)—a puzzle that proponents of the Faultlessness View avoid.

For these reasons, I am inclined to think that the Faultlessness View is the most promising version of a deontological approach. But regardless of which version of a deontological approach we settle upon, it should be clear that a deontological approach can be combined with an attitudinal metaepistemology to deliver a reductive semantics for justification ascriptions. According to each version of a deontological approach, *justified* quantifies over a set of accessible worlds ranked by an ordering source. And according to the attitudinal metaepistemology developed in §2.3.1, the ordering source is provided by certain agents’ states of epistemic approval.

²⁷Harman (2016) calls cases along these lines, ‘morally permissible mistakes.’

How exactly does a deontological approach (conjoined with an attitudinal metaepistemology) handle cases of defeat? According to the approach developed here, the semantics specifies truth conditions for justification ascriptions in terms of agents' states of epistemic approval. But it does not tell us *what* the agents in question epistemically approve of. This is a question to be settled by the psychology, not the semantics. It is thus in an important sense a 'black box', one that we can help ourselves to when applying the semantic framework.

To illustrate, let us walk through Consuela and the Vase again. Consider the initial time t_1 when Consuela first sees a red vase in front of her (prior to receiving her interlocutor's testimony). Now suppose that the agent supplied by the context of utterance is just *you*, dear reader. Now ask yourself: Do you epistemically approve of Consuela's belief that the vase in front of her is red? If you're anything like me, the answer is 'Yes'. And so, according to the deontological approach, (20) ('Consuela is justified in believing that there's a red vase in front of her') will be true at t_1 .²⁸

Now consider the later time t_2 , after Consuela has received her interlocutor's testimony. Do you still epistemically approve of Consuela's belief that the vase in front of her is red? Presumably, the answer is 'No'. And so the deontological approach predicts that (20) is false at t_2 .

While I have focused on epistemic justification ascriptions, the deontological approach can be easily extended to handle moral justification ascriptions, as in:

(29) Sasha was justified in injuring her assailant.²⁹

For example, proponents of the Faultlessness View will analyze (29) as saying that it's not the case that Sasha (morally) shouldn't have injured her assailant. And for proponents of an attitudinal metaethics, this deontic claim can in turn be analyzed in terms of existential

²⁸Of course, the different versions of the deontological approach will yield different predictions depending on exactly which worlds are consistent with your epistemic approval profile. As long as every world that is acceptable, from the point of view of your epistemic approval profile, is a world where Consuela holds the belief in question, then all versions of the deontological approach will predict that (50) is true at t_1 . By contrast, if Consuela only holds this belief in some of the epistemically acceptable worlds but none of the epistemically optimal worlds (again, from the point of view of your epistemic approval profile), then only the Permissive View predicts that (20) is true at t_1 .

²⁹Corpus searches suggest that moral uses of *justified* are far more common in ordinary English than epistemic uses.

quantification over the worlds in the modal base that are ranked highest by certain agents' moral approval profiles.

A Scalar Semantics for *Justified*

A different strategy for analyzing *justified* starts from the observation that *justified* is a gradable adjective. The two general diagnostics for whether an adjective *A* is gradable is whether it admits of degree modifiers (*very*, *fairly*, etc.) and whether it admits of comparatives. The term *justified* passes both of these diagnostics with flying colors:

- (30) The jurors are very justified in believing the defendant is guilty.
- (31) Given the evidence, we're justified in believing that all of the search committee members are corrupt. But we're even more justified in believing that at least one of them is corrupt.

More specifically, *justified* seems to function as a *relative gradable adjective*—an adjective that is associated with some contextually determined threshold.³⁰ The main test for whether *A* is a relative gradable adjective is whether we can coherently say things like, '*x* is *A*, but it could be more (/less) *A*.' For example, *tall* is a relative gradable adjective, as evidenced by the felicity of the following:

- (32) Sarah is tall, but she could be even taller (if, for example, she grew a couple of inches).
- (33) Sarah is tall, but she could get shorter (if, for example, she shrunk a couple of inches).

The term *justified* seems to work similarly:

- (34) Poirot is justified in believing the butler did it, but he could be even more justified in this belief (if, for example, he gathered more evidence that the butler is guilty).
- (35) Poirot is justified in believing the butler did it, but he could be less justified in this belief (if, for example, he acquired a defeater).

³⁰Relative adjectives are typically contrasted with *absolute gradable adjectives*, which will be discussed in Chp.3.

In the semantics literature, the most popular analysis of gradable adjectives is a scalar semantics.³¹ On this approach, every gradable adjective A is associated with a scale Θ_A , where a scale is a set of degrees ordered along some dimension. In the case of *tall*, the scale might be height; in the case of *expensive*, the scale will be cost, etc. According to the standard approach, an ascription of a relative gradable adjective of the form, ‘ x is A ’ will be true if and only if x ’s degree on the A -associated scale Θ_A exceeds some contextually determined threshold.

To illustrate, consider again *tall*. Let Θ_{tall} be the scale associated with *tall* (presumably, a scale consisting of various heights). Let **tall** be a measure function that, for any object x , yields x ’s value on Θ_{tall} (that is, x ’s height). Then we can offer the following semantics for *tallness* ascriptions:

$$\llbracket x \text{ is tall} \rrbracket^{c,w} = 1 \text{ iff } \mathbf{tall}(x) > d$$

(where d is some c -determined threshold)

One way of developing a scalar semantics for *justified* in the context of an attitudinal metaepistemology is to let the underlying scale be *degrees of epistemic approval*. Here’s a simple contextualist implementation of this analysis. Let \mathbf{E}_A^w be a measure function that, for any proposition p , delivers the degree to which A epistemically approves of p at w . Then we can offer the following semantics for justification ascriptions:

$$\llbracket S \text{ is (epistemically) justified in } \phi\text{-ing} \rrbracket^{c,w} = 1 \text{ iff } \mathbf{E}_A^{w_c}(\{w': \llbracket S \phi's \rrbracket^{c,w'} = 1\}) > d$$

(where A is some c -determined subject, and d is some c -determined threshold).

Now consider the special case of this where the value of ‘ ϕ -ing’ is ‘believing p .’ We get the following analysis:

$$\llbracket S \text{ is justified in believing } p \rrbracket^{c,w} = 1 \text{ iff } \mathbf{E}_A^{w_c}(\{w': S \text{ believes } p \text{ at } w'\}) > d.$$

According to this semantics, a sentence such as (20) (‘Consuela is justified in believing there’s a red vase in front of her’) will be true in a context c iff the degree to which some

³¹See Cresswell 1977; Kennedy and McNally 2005; Kennedy 2007; Lassiter 2011, among many others.

contextually determined subject epistemically approves of Consuela believing there's a red vase in front of her exceeds the contextually determined threshold. (As in the case of modals, we could also tweak this proposal in various ways—for example, by putting a relativist or an expressivist spin on it.)³²

How does this approach handle cases of defeat? In much the same manner as the deontological approach. Suppose again that the contextually determined agent is just you, the reader. Presumably, at t_1 you epistemically approve of Consuela believing that there's a red vase in front of her, and you do so to a sufficiently high degree. Presumably, at t_2 you do not. Given this, the scalar approach agrees with the deontological approach: (20) is true at t_1 and false at t_2 .

2.3.3 Beyond Modals and Justification Ascriptions

So far I've shown how one can develop an attitudinal semantics for epistemic evaluations involving deontic modals (§2.3.1) and justification ascriptions (§2.3.2). The present approach can also be extended to other epistemic appraisals.

Take, for example, knowledge ascriptions. Suppose we adopted the toy view that knowledge is justified safe belief. Then, we would wind up with the view that an ascription of the form, 'S know p ' is true in a context c iff (i) S couldn't easily have falsely believed p , (ii) 'S's belief that p is justified' is also true in c . Then we could use our reductive semantics for justification ascriptions to give a reductive treatment of knowledge ascriptions.

Of course, a simple analysis of knowledge as justified safe belief may well be mistaken. Indeed, it may turn out that knowledge is unanalyzable (Williamson 2000a).³³ Still, it seems plausible that knowledge entails justification, i.e.:

- (36) a. S know p . \Rightarrow
 b. S is (epistemically) justified in believing p .

And if this is right, then at the very least we can use our reductive semantics for justification ascriptions to give a reductive account of one important facet of knowledge.

³²In Chp.3 I develop the view that degrees of justification are best understood in terms of degrees of epistemic certainty. This proposal could be rendered extensionally equivalent to the current approach if we analyze degrees of epistemic certainty in terms of degrees of epistemic approval.

³³I discuss this issue in more detail in chp.3, where I suggest an account of knowledge as near certainty.

Similar remarks apply to ascriptions of evidence possession (that is, sentences of the form: ‘S has p as evidence’). If we understand evidence as justified belief, or as justified safe belief, etc., then we can leverage our reductive semantics for justification ascriptions into a reductive semantics for evidence possession ascriptions. And even if such a simple analysis turns out to be mistaken, then as long as we accept the Evidence-Justification Link (Chp. 1), we can use our reductive semantics for justification ascriptions to get a partial purchase on evidence possession ascriptions.

2.4 Advantages of an Attitudinal Metaepistemology

From the present perspective, the main advantage of an attitudinal metaepistemology is its reductive payoff: we can give semantic analyses of a variety of epistemic linguistic expressions in terms of conative attitudes, which themselves can be understood without reliance on any unreduced epistemic notions. Moreover, as we’ve seen, the basic strategy here nicely parallels the strategy behind an attitudinal metaethics: we reduce the normative to the non-normative by way of the psychological.

But there are additional benefits to be had. In particular, an attitudinal metaepistemology, when conjoined with an attitudinal metaethics, illuminates many of the close parallels between epistemic and moral discourse. It’s to these parallels that I now turn.

2.4.1 Value-Ladenness

Both moral and epistemic evaluations appear to be *value-laden*.³⁴ If I say that the U.S. (morally) ought to allocate more money to foreign aid, it seems my claim commits me to valuing a state of affairs in which the U.S. allocates more money to foreign aid. Similarly, if I claim that we (epistemically) ought to believe in the external world, it seems my claim commits me to valuing a state of affairs in which we believe in the external world. In this regard, moral and epistemic evaluation seem to differ from purely descriptive claims (‘Grass is green’ and the like) which don’t seem to commit speakers to valuing any particular state

³⁴ Arguably, this is equivalent to saying that moral and epistemic evaluations are both *normative*. However, due to the controversy over how to define normativity, I’ll use the less theoretically fraught locution ‘value-laden’ instead.

of affairs.

Proponents of an attitudinal semantics have a straightforward explanation of the value-ladenness of epistemic and moral evaluations. The explanation goes like this: moral and epistemic approval are both evaluative attitudes; to morally approve of something is to value it in a certain way, to epistemically approve of something is to value it in a different way. Moral and epistemic evaluations count as value-laden in virtue of the evaluative nature of the moral and epistemic attitudes that are intimately involved in their truth-conditions.

One attractive feature of this explanation of the value-ladenness of moral and epistemic evaluations is that it extends to other value-laden evaluations. Consider, for example, prudential evaluations such as:

- (37) Given her interests, she ought to take Chemistry.

The natural way of analyzing prudential evaluations in the current framework is to posit a further conative attitude: approval of some state of affairs insofar as that state of affairs is conducive to a particular agent's interests. Arguably, we've already looked at some examples of this attitude: Shelly approves mediocre Mel's belief that he's a talented musician insofar as this belief is conducive to Mel's interests, and the avaricious business mogul approves of hoarding his wealth insofar as this is conducive to his own interests. We can then use this attitude to provide the ordering source for prudential uses of *ought*. On the picture that emerges, all value-laden discourse is analyzed in terms of certain conative attitudes. Different flavors of value-ladenness (moral value vs. epistemic value vs. prudential value) are explained in terms of differences between these conative attitudes.

2.4.2 Connection with Motivation

Many metaethicists have been attracted to the view that moral judgments are intimately bound up with motivation.³⁵ To see the appeal of this thesis, imagine encountering someone who insists that everyone is morally obligated to give to charity. We'd naturally expect her

³⁵Moral judgment internalists hold that there is a necessary connection between moral judgment and motivation: see Dreier 1990; Smith 1994; Korsgaard 1996; Blackburn 1998; Gibbard 2003; Wedgwood 2007; van Roojen 2010; Egan 2012, among many others. Even moral judgment externalists typically grant that there is a close connection between moral judgment and motivation; they just deny that this is a necessary connection.

to be at least somewhat disposed to give to charity.

Of course, we should be careful not to exaggerate the connection between moral judgment and motivation; as our avaricious business mogul shows, akrasia is a familiar phenomenon.³⁶ In order to avoid counterexample, our statement of the connection between moral judgment and motivation had better be suitably qualified.

The literature offers a variety of proposals for how to formulate a suitably qualified connection between moral judgment and motivation. For my purposes, I will operate with the following fairly weak statement of the connection:

Moral Judgment-Motivation Connection: Typically, if a psychologically normal agent believes that she's morally obligated to ϕ , then she'll be at least somewhat disposed to ϕ .^{37,38}

(The details of this formulation won't matter for my purposes; if you prefer a somewhat different formulation of the connection between moral belief and motivation, most of what I say in the following should still go through.)

Just as there is pressure to accept some connection between moral judgment and motivation, there's pressure to accept an analogous connection in the epistemic domain. In the epistemic domain, the connection seems not to be between epistemic judgments and actions, but rather between epistemic judgments and doxastic attitudes. We already observed (§2.2) that it would be rather surprising to affirm that one is epistemically obligated to adopt some doxastic attitude D towards p , while exhibiting no inclination whatsoever to adopt D towards p . Here is one way of formulating such a connection between epistemic judgment and motivation:

Epistemic Judgment-Motivation Connection: Typically, if a psychologically normal agent sincerely judges that she's epistemically obligated to adopt doxastic attitude

³⁶For discussion of the difficulties that akratic subjects pose for strong versions of internalism, see Stocker 1979; Smith 1994; Mele 1996; Svavarsdóttir 1999.

³⁷For formulations of internalism in terms of what's 'psychologically normal,' see Dreier 1990; Blackburn 1998, among others.

³⁸Note that since this statement of the connection is formulated in terms of what *typically* holds (rather than what *necessarily* holds), it's a thesis that even motivational externalists can get on board with.

D towards p , then she'll be at least somewhat disposed to adopt D towards p .³⁹

In short, it's plausible that there's at least some connection between one's moral judgments and one's motivation to act—a connection of the sort articulated by Moral Judgment-Motivation Connection. However, it's equally plausible that there's a similar connection between epistemic evaluations and one's motivation to believe—a connection of the sort described by Epistemic Judgment-Motivation Connection.

Many philosophers hope to provide a metaethical explanation of the connection between moral judgment and motivation.⁴⁰ It seems just as natural to hope for a metaepistemological explanation of the connection between epistemic judgment and motivation. Perhaps if an attitudinal metaethics can underwrite the Moral Judgment-Motivation Connection, an attitudinal metaepistemology can underwrite the Epistemic Judgment-Motivation Connection. The basic idea would be this: just as we can appeal to the motivational efficacy of moral approval to explain the motivational pull of moral judgments, we can appeal to the motivational efficacy of epistemic approval to explain the motivational pull of epistemic judgments.

This strikes me as a promising explanatory strategy. But how exactly should we spell out the details? Expressivism offers what is perhaps the most well-known way of using an attitudinal metaethics to explain the Moral Judgment-Motivation connection. Expressivists typically embrace the idea that moral beliefs are *first-order states*: to believe that p (morally) ought to be the case is *not* to believe that certain agents morally approve of p , or would morally approve of p in such-and-such conditions. Rather, it's just to stand in a desire-like relation (e.g., moral approval) towards p itself. Thus a sentence such as:

(38) Jane believes that everyone (morally) ought to give to charity.

does not say that Jane believes some proposition of the form: ⟨I morally approve of everyone giving to charity⟩. Rather, it ascribes to Jane a first-order state of mind: it says that Jane morally approves of everyone giving to charity.

³⁹Some might be wary of any talk of 'motivation' to believe, since motivation seems intimately bound up with volition, and it's highly controversial whether we have volitional control over our beliefs. To allay this worry, I should stress that nothing hinges on my use of 'motivation' talk; one should feel free to regard this talk as a convenient metaphor for dispositions to belief.

⁴⁰See Stevenson 1963; Blackburn 1984; Smith 1994; Dreier 1990; Egan 2012, among others.

Extended to epistemic expressivism, the idea would be that epistemic judgments are also first-order states: all it is to have an epistemic belief is to be in a desire-like state (e.g., epistemic approval). For example, according to epistemic expressivism:

(39) Jane believes that everyone (epistemically) ought to believe in the external world.

does not say that Jane believes that she epistemically approves of everyone believing in the external world. Rather, it just says that Jane epistemically approves of everyone believing in the external world.

Here's one way of semantically implementing this idea within the present framework. Start with the familiar idea that *believes* quantifies over a subject's doxastic alternatives, where a subject S's doxastic alternatives at a world w (Dox_S^w) are the worlds compatible with what S believes at w (Hintikka 1962). However, we add a tweak: the verb *believes* also shifts the ordering source, which—on the current framework—is provided directly by moral and epistemic approval states in the index. In particular, it shifts these parameters so that the ordering is provided directly by the believer's states of moral and epistemic approval, i.e.:

Expressivist Semantics for *Believes*:

$$\llbracket S \text{ believes } \phi \rrbracket^{c,f,w,e,m} = 1 \text{ iff } \forall w' \in Dox_S^w : \llbracket \phi \rrbracket^{c,f,w',e_S^w,m_S^w} = 1.$$

By combining this with the Extended Expressivist Semantics for *Ought*, we get the desired predictions. (38) ascribes to Jane a first-order state of mind: roughly, the state of morally approving of everyone giving to charity.⁴¹ And (39) also ascribes to Jane a first-order state of mind: roughly, the state of epistemically approving of everyone believing in the external world.

Equipped with this semantics for belief reports, we can give a unified expressivist explanation of the motivational efficacy of moral and epistemic judgments. To see this, let's start with epistemic judgments. As we've already seen, it's arguably part of the functional role of epistemic approval that, *ceteris paribus*, a psychologically normal agent who believes

⁴¹More precisely, the Expressivist Semantics for *Believes* says that to compute the truth conditions for (38), we take each of Jane's doxastic alternatives w' and plug it into the modal base function, getting a set of worlds ($f(w')$). If for every such set of worlds, the worlds that are ranked highest by Jane's moral approval profile are worlds where everyone gives to charity, then (38) is true.

that she ought to have doxastic attitude D towards p will be disposed to adopt D towards p . And so if it is indeed true that Jane epistemically approves of everyone believing in the external world, then, *ceteris paribus*, Jane will—if psychologically normal—be disposed to believe in the external world.

The same basic story can be told in the moral domain. Plausibly, it's part of the functional role of moral approval that psychologically normal agents who morally approve of ϕ -ing in circumstances σ will be disposed to try to ϕ if they take themselves to be in σ . And so if it is indeed true that Jane morally approves of giving to charity, then, *ceteris paribus*, Jane will—if psychologically normal—be disposed to give to charity.

Thus, by combining an expressivist metaepistemology with an expressivist metaethics, we can give the unified explanation that we sought. Can we run a similar explanation in the contextualist or relativist framework? Typically, contextualists and relativists have denied that normative beliefs are first-order states of mind. For example, Egan (2012), building on Lewis (1989), takes value judgments to be *de se* attitudes: specifically, self-ascriptions of properties such as: *being disposed to undergo certain responses in certain circumstances* (e.g., if idealized in certain ways).

Here's a simple semantic implementation of this idea. Suppose we follow Lewis (1979a) and take doxastic alternatives to be centered worlds (where, as before, a centered world is an ordered pair of a world and some agent). As before, we say that *believes* universally quantifies over the believer's doxastic alternatives. However, we add that *believes* shifts the assessor parameter, so that the complement clause is evaluated relative to the center of each of these alternatives:

Relativist Semantics for *Believes* (First Pass):

$$\llbracket S \text{ believes } \phi \rrbracket^{c,f,g,w,a} = 1 \text{ iff } \forall \langle w', i \rangle \in \text{Dox}_S^w : \llbracket \phi \rrbracket^{c,f,g,w',i} = 1,$$

where $\text{Dox}_S^w = \{ \langle w', i \rangle : \text{it's compatible with what } S \text{ believes at } w \text{ that } S \text{ is } i \text{ at } w' \}$.⁴²

By combining this semantics for *believes* with the Relativist Semantics for *Ought*, we get what Egan takes to be the right result: (38) ('Jane believes that everyone (epistemically)

⁴²Here I draw on Stephenson (2007a,c), who offers a similar treatment of attitude reports embedding taste predicates and epistemic modals.

ought to believe in the external world') says that Jane self-ascribes the property of *being someone who epistemically approves of everyone believing in the external world* (or the property of *being someone who, if idealized in ways x , y , and z , would epistemically approve of everyone believing in the external world*, etc.). Similarly, (39) ('Jane believes that everyone (morally) ought to give to charity') says that Jane self-ascribes a certain property, e.g., the property of *being someone who morally approves of everyone giving to charity*.

Does this semantics for *believes* underwrite the connection between the normative judgment and motivation? This is less clear. Unlike the expressivist semantics for *believes*, our first pass relativist approach does not treat normative belief as a first order state: to hold a normative belief is not to be in a desire-like state, but rather to have a *de se* belief that one is in such a state. Thus our first pass relativist approach allows for a distinction between the *good case*, in which one is in the desire-like state in question, and the *bad case*, in which one isn't. Underwriting the connection with motivation in the good case is easy: if Jane does indeed instantiate the property of *being someone who epistemically approves of everyone believing in the external world*, then we can once again appeal to the functional role of epistemic approval to explain why we'd expect Jane to be at least somewhat disposed to believe in the external world. By contrast, underwriting the connection with motivation in the bad case is more difficult: if Jane isn't someone who epistemically approves of everyone believing in the external world, then we can't appeal to the functional role of epistemic approval to explain why her normative belief seems to carry motivational pull.

How worrisome should we find this lacuna? This isn't entirely clear to me. There are various ways that proponents of an Egan-style view might try to handle the bad case. First, they might point out that we typically presume that people have fairly reliable (albeit fallible) introspective access to their own conative attitudes. And so perhaps we have a presumption that the good case is the default: if Jane self-ascribes the property of being in a certain state of epistemic or moral approval, we'll typically presume that she is in the state in question. Alternatively, they might appeal to the idea that most people have higher-order desires to act in accordance with their states of moral or epistemic approval. (Most people want to be enkratic.) And so perhaps if Jane mistakenly self-ascribes the property of morally approving of giving to charity, we would expect this self-ascription to

combine with her desire to act in accordance with her states of moral approval in order to generate some disposition to give to charity.

Adjudicating whether these strategies suffice to underwrite the connection with motivation is a difficult matter. Much here depends on exactly how strong a connection between normative judgment and motivation one endorses: the stronger the connection, the more difficult it will be for our first pass relativist semantics for *believes* to underwrite it. For those who are pessimistic about the relativist's prospects here, it is worth noting that relativists are not wedded to an Egan-style view on which normative beliefs are second-order states. Instead, they could tweak their semantics for *believes* so as to deliver the result that normative beliefs are first-order states.

Here's a simple way to do this. As in the expressivist framework, let an agent's doxastic alternatives just be worlds (rather than centered worlds). We then propose that *believes* simply shifts the value of the assessor parameter to the believer, i.e.:

Relativist Semantics for *Believes* (Second Pass):

$$\llbracket S \text{ believes } \phi \rrbracket^{c,f,g,w,a} = 1 \text{ iff } \forall w' \in \text{Dox}_S^w : \llbracket \phi \rrbracket^{c,f,g,w',S} = 1.$$

This approach delivers the same predictions as the expressivist approach when it comes to normative belief reports. In particular, it predicts that (38) and (39) ascribe to Jane first-order states of moral and epistemic approval.

There is an interesting moral to be extracted from this. In the metaethics literature, most have assumed a tight connection between expressivism (construed as either a semantic or a pragmatic thesis) and the thesis that moral beliefs are first-order states. However, the latter thesis is in principle compatible with semantic and pragmatic theses that, at least historically, have been considered rivals to expressivism: namely, relativism and contextualism.⁴³

⁴³Here's one way of getting the result that normative beliefs are first-order states within a contextualist framework. Let g_S^w be a constant function that takes worlds as inputs and yields S's epistemic or moral approval profile at w . (Which it delivers will depend on the flavor of deontic modality in question, which will be determined by the context.) The contextualist can now propose the following semantics for *believes*:

$$\llbracket S \text{ believes } \phi \rrbracket^{c,f,g,w} = 1 \text{ iff } \forall w' \in \text{Dox}_S^w : \llbracket \phi \rrbracket^{c,f,g_S^w,w'} = 1.$$

This also delivers the same result as our Expressivist Semantics for *Believes* and our second pass relativist semantics: moral and epistemic beliefs are just first-order states.

Let's take stock. It's plausible that there is some close connection between moral belief and motivation to act, and that there's a similarly close connection between epistemic judgment and motivation to believe/disbelieve. (Though exactly how close of a connection is open to dispute.) By adopting an attitudinal metaethics and metaepistemology, we can give a unified explanation of these two connections: both are explained in terms of the motivational pull of the underlying psychological states. Exactly how the explanation goes will depend on one's semantics for belief reports (and, more generally, whether one takes ethical and epistemic beliefs to be first-order or second-order states). While I have not tried to settle the question of what is the *right* attitudinal semantics for belief reports, I hope to have shown that there are some promising options on the table, options that—when combined with our expressivist, relativist, or contextualist semantics (§2.3)—ensure that the desired connections between normative belief and motivation hold.

2.5 Comparison with a First-Order Reduction

I've been arguing that an attitudinal metaepistemology can deliver what we failed to find in the first chapter: an account of how the epistemic reduces to the non-epistemic. But some may question whether it really delivers that much.

Here's one way of developing the worry. Suppose that someone utters an epistemic evaluation such as (30):

(30) The jurors are justified in believing the defendant is guilty.

Suppose moreover that this epistemic evaluation is true (relative to whatever the relevant parameters are). Now a question arises:

Grounding Question: What makes this epistemic evaluation true?

Presumably, it's wrong to say that this evaluation is true *in virtue of* certain agents' states of epistemic approval. Or at least, even if it's true partially in virtue of such states, that's not the full story: presumably there's some further explanation to be given.

In trying to give this further explanation, the natural thing to do is find some further facts $f_1...f_n$ about the jurors' situation in virtue of which the relevant agents epistemically

approve of the jurors believing that the defendant is guilty. Now, our objector may observe, the attitudinal framework I've developed gives us no guidance as to what these further facts are. Indeed, for all I've said, $f_1...f_n$ could correspond to any of the first-order reductions discussed in the previous chapter: they could be facts about the reliability of the jurors' belief-forming processes, or facts about jurors' non-factive mental states, etc. A fully satisfactory reductive account would give us some general formula for identifying such facts. Only once we have such a formula will we have made good claim to vindicating Reducibility.

It is certainly true that I haven't given any recipe for identifying which facts ground various agents' states of epistemic approval or disapproval. But perhaps it is too much to ask for such a recipe. Plausibly, agents differ in terms of what facts elicit their states of epistemic approval/disapproval (as evidenced by conflicts in people's epistemic intuitions). And even if we hone in on a single agent (for example, you, the reader), it may well prove difficult or impossible to specify, in non-epistemic terms, which facts about a situation elicit their states of epistemic approval/disapproval. Indeed, this may be precisely the lesson we should draw from our failure—in the first chapter—to come up with a counterexample-free first-order reductive epistemology: if we're unable to specify in non-epistemic terms the conditions under which a subject S is justified in believing p , we'll likewise be unable to specify in non-epistemic terms the conditions under which we epistemically approve of S believing p .

This suggests the following possibility: whenever we do epistemically approve or disapprove of some particular subject adopting some doxastic attitude towards some proposition, our epistemic approval or disapproval is responsive to certain non-epistemic facts about that subject's situation. Still, it may be that any general specification of the non-epistemic facts that are necessary and sufficient to elicit our epistemic approval of a subject's belief will be highly disjunctive, and hence that all first-order reductions are bound to fail. To put it in terms drawn from the moral particularism literature, it might be that grounds for our states of epistemic approval are “uncodifiable.” While this possibility would be somewhat disappointing, the attitudinal framework developed here offers some consolation. If the attitudinal framework developed here is right, we can still give a high-level, non-disjunctive account of the relation between the epistemic and the non-epistemic. It's just that the

account in question doesn't take the form of a specification of the grounds for epistemic facts; rather, it takes the form of a semantic theory mapping epistemic sentences to their meanings, where these meanings are characterized in non-epistemic terms.⁴⁴

This possibility stands in stark contrast to Goldman's opening remarks on epistemic reduction in his classic paper, "What is Justified Belief?". There he writes:

The term 'justified', I presume, is an evaluative term, a term of appraisal. Any correct definition or synonym of it would also feature evaluative terms. I assume that such definitions or synonyms might be given, but I am not interested in them. I want a set of *substantive* conditions that specify when a belief is justified. Compare the moral term, 'right'. This might be defined in other ethical terms or phrases, a task appropriate to metaethics. The task of normative ethics, by contrast, is to state substantive conditions that determine whether an action is right. (Goldman 1979: 90)

In this passage, Goldman suggests that any correct semantic analysis of epistemic expressions such as *justified* would have to invoke some normative notions, and hence fail to be reductive. However, Goldman is considerably more optimistic about the prospects for a first-order reduction. According to the possibility that I am suggesting, the situation is precisely the opposite: we can give a semantic reduction (of the attitudinal variety); but the prospects for a first-order reduction are bleak.

Now, it may turn out that my pessimism about first-order reductions is premature. Perhaps with time we'll be able to develop a sufficiently refined version of ARP that avoids

⁴⁴Some may question the extent to which this is a genuine consolation. After all, a semantic reduction seems to be a very different sort of reduction than that offered by a first-order account. In what sense, then, does it answer the same questions that prompted the search for the latter?

Admittedly, it is rather difficult to say precisely how a semantic reduction relates to a first-order reduction. But here is a stab at spelling out the connection. One motivation for some sort of reductive analysis is the desire to understand how the epistemic relates to the non-epistemic. One face of the epistemic is linguistic (and, more generally, conceptual): it consists in the terms and concepts we use in making epistemic appraisals. And so one way of relating the epistemic to the non-epistemic is to show how epistemic terms and concepts can themselves be understood in non-epistemic terms. (Here too, it is worth stressing the comparison with attitudinal accounts in metaethics. For example, Blackburn (1998) and Gibbard (2008) both take themselves to be explaining the relation between the normative and the natural, the place of *oughts* in the world of *iss.*) Another face of the epistemic is metaphysical: it consists in the properties and states denoted by our epistemic terms and concepts. And so another way of relating the epistemic to the non-epistemic is to show how epistemic properties and states metaphysically reduce to non-epistemic properties and states. What unites the two reductive projects, according to this proposal, is that both are prompted by a desire to understand the relation between the epistemic and the non-epistemic.

the counterexamples in chp.1, or a reductive account of evidence possession that validates the Evidence-Justification Link. If so, then such an account could be combined with the attitudinal framework in a way that answers the Grounding Question: the right first-order account (whatever that is) could be used to provide a general formula for specifying which non-epistemic facts ground the relevant agents' states of epistemic approval/disapproval.

But if some such first-order reduction were forthcoming, wouldn't this undermine the motivation for an attitudinal metaepistemology? The answer is: "Only to some extent." We would no longer need to give a reductive metaepistemology as a consolation prize for failing to have arrived at a first-order reductive account. However, a reductive metaepistemology seems worthwhile in its own right: if, on top of being able to specify the grounds of epistemic facts in non-epistemic terms, we can also specify the meanings of epistemic expressions in non-epistemic terms, so much the better. What's more, I have argued that the attitudinal framework holds attractions above and beyond its reductive *bona fides* (§2.4). In particular, it sheds light on some of the commonalities between epistemic and moral language, providing a unified explanation of both (i) the value-ladenness of epistemic and moral expressions, and (ii) their shared connection with motivation. This is an advantage that should make the attitudinal framework appealing even to those who are confident in the prospects of a first-order reductive account, and even for those who are uninterested in reduction altogether.

Chapter 3

The Centrality of Certainty

3.1 Introduction

For much of its history, epistemology focused on certainty. In the medieval and early modern periods, epistemological discussions centered on *scientia*, conceived as the highest grade of knowledge. Most authors in this tradition explicitly associated *scientia* with certainty. For example, we find Aquinas, Scotus, and Descartes all characterizing *scientia* in terms of ‘certain [*certa*] cognition.’¹

For these authors, *scientia*—and more specifically, certainty—is the optimal epistemic status: it is ‘perfect cognition.’ And while certainty was conceived as the epistemic ideal, it was deemed an attainable ideal. This epistemic optimism is perhaps clearest in Descartes, who urged that by properly directing our reason, we can elevate much of our everyday knowledge to the status of *scientia*.

But as time went on, this optimism waned. According to an emerging consensus, most of our knowledge only approaches the optimal epistemic status; precious little actually reaches this ideal. Thus we find Russell (1912) concluding that only our knowledge of our sensory states makes the cut—all else is merely probable. In a similar vein, Reichenbach concludes that all attempts to achieve certainty have failed: “The search for certainty had to burn itself out” (1963: 49). A sophisticated development of this line of thought is found in Unger

¹Aquinas 1970; Scotus 1987, V; Descartes 1983, X: 362. See Pasnau (2014a,b) for detailed discussion of the connection between *scientia* and certainty in the medieval and early modern tradition.

1971, 1975, who argues that *certain* functions as an absolute term, which only applies to a proposition if that proposition has the maximal degree of certainty. This is then conjoined with a general argument that absolute terms almost never apply to anything in the world.

If certainty is so rarely attained, it is natural to wonder whether it should matter much for epistemology. For many contemporary epistemologists, the answer seems to be ‘No.’ According to a common line of thought, while certainty is largely beyond our reach, knowledge is much more easily attained; consequently, it’s better-suited to play a central role in epistemology. This shift in focus from certainty to knowledge is perhaps most apparent in the ‘Knowledge First’ program in epistemology. For Knowledge Firsters, knowledge is the central epistemic notion: it can be used to illuminate other epistemic notions (e.g., evidence, evidentiality probability, justification) and to provide normative constraints on assertion, belief, and action.² Typically, Knowledge Firsters devote scant space to certainty; certainly the notion doesn’t play an important explanatory role in extant knowledge-centric accounts.

In this chapter, I seek to restore certainty to its former centrality. I open by inquiring into the nature of certainty. I argue that attempts to reduce certainty to knowledge flounder: certainty is a *sui generis* epistemic notion. At the same time, I contest the perception of certainty as an unattainable ideal. I argue that much—though by no means all—of our ordinary knowledge qualifies, in appropriate contexts, as certain.

Having laid the groundwork, the rest of the chapter explores how much explanatory mileage we can get out of certainty: if we were to take certainty as our primitive, how far could this take us? Surprisingly far, it turns out. Certainty can be used to provide illuminating accounts of a variety of other epistemic notions, including evidential probability, epistemic modals, and knowledge. It also provides plausible normative constraints on credence, assertion, and action. I suggest that these results warrant cautious optimism about the prospects of a ‘Certainty First’ epistemology. At the very least, they show that certainty has a relatively high degree of explanatory priority in the epistemic realm, and hence should play a central role in epistemology.

²See Williamson 1996, 1997, 1998, 2000a,b, 2011, 2013; Hawthorne 2004; Hawthorne and Stanley 2008; Stanley 2005a; Sutton 2005, 2007; Smithies 2012; Littlejohn forthcoming.

3.2 Certainty: A Schematic Account

3.2.1 Subjective vs. Epistemic Certainty

It's common to distinguish between subjective and epistemic certainty.³ Subjective certainty is a matter of strength of conviction. A belief can be subjectively certain even if it's held for no good reason. By contrast, if a belief is epistemically certain, the believer must stand in a strong epistemic relation to its content.

While *certain* and its cognates are ambiguous between these two senses, certain constructions favor one reading over the other (Moore 1959; Stanley 2008; DeRose 2009). Claiming that a person is certain of something usually conveys subjective certainty:

(40) I'm certain/sure that the butler did it.⁴

(40) can be true even if the speaker falsely and irrationally believes the butler did it.

Claiming that a proposition is certain usually conveys epistemic certainty:

(41) It's certain that the butler did it.

(41) seems to entail that the speaker is in a strong epistemic position with regards to the proposition that the butler did it.

What sort of epistemic position? Some might suggest it's knowledge. But this suggestion has trouble explaining the fact that *knows for certain* is not redundant. To see this, imagine that it's the first day of Epistemology 101, and you're trying to get your students to feel the pull of Descartes' project. Most likely, you'd ask (42a) rather than (42b):

(42) a. What can we know for certain/with certainty?

b. What can we know?⁵

³See e.g. Moore 1959; Klein 1981; Stanley 2008; DeRose 2009; Reed 2011.

⁴I'll assume that *certain* and *sure* are roughly synonymous, despite subtle differences in their distributions. For a corpus-based discussion of such differences, see Simon-Vandenbergen and Aijmer 2007.

⁵Arguably, asking (42a) rather than (42b) fits better with Descartes' own views on knowledge. While Descartes is widely interpreted as holding that knowledge requires certainty, Descartes' discussion of the atheist mathematician in the Second Replies casts doubt on this interpretation. In his discussion, Descartes draws a distinction between *cognitio* and *scientia*: the atheist's belief that three angles of a triangle are equal to two right angles amounts to *cognitio*, but not *scientia* (AT VII 141). On a natural reading, *cognitio* still amounts to a species of knowledge, it is simply a lower grade than *scientia*. For further discussion, see Sosa 1997; Wykstra 2008; Pasnau 2014a.

For another illustration of the non-redundancy of *knows for certain*, note that (43a) seems to make a stronger claim than (43b):

- (43) a. Alex knows for certain/sure that Jean was at the store last night.
 b. Alex knows that Jean was at the store last night.

One might try to account for the non-redundancy of *knows for certain* by insisting that whenever *certain* combines with *knows*, it conveys subjective rather than epistemic certainty. But this seems wrong. It seems that there's a very natural reading of (43a) on which the speaker isn't claiming that Alex knows Jean was at the store and that, furthermore, Alex is extremely confident of this fact. Rather, the speaker is claiming that Alex knows this fact with a particularly high degree of justification or warrant.⁶

For further evidence that epistemic certainty is a more demanding state than knowledge, consider the following examples from the wild:

- (44) [W]e know without certainty, but with a high degree of probability, that returns over the next 10 years or so will be very poor.⁷
 (45) When [a false ID] is handed to a cop, he knows with near certainty the guy before him is not the guy identified on the flimsy piece of paper.⁸

If epistemic certainty were the same as knowledge, such claims would be incoherent.

⁶Some might try to resist this argument by maintaining that both *certain* and *knows* are context-sensitive expressions, and that combining them in a complex phrase such as, *knows for certain* drives up the standards for both knowledge and certainty. However, it proves difficult to produce a plausible general principle to explain why this would be the case. In general, if two expressions are broadly synonymous, preposition or conjunction phrases combining them are not naturally interpreted as invoking heightened standards. For example, saying that the detective is *sure and certain* sounds redundant; it's not naturally interpreted as saying that the detective is certain to an unusually high degree. (For further illustration, consider how we would interpret the claim that something is *likely and probable*, or that someone *understands with comprehension*.)

Alternatively, one might maintain that *knows for certain* is an idiom, and consequently that its meaning is not determined compositionally from the meanings of *knows* and *certain*. One difficulty with this hypothesis is that idioms tend not to be cross-linguistically widespread: in English, *kick the bucket* is an idiom for dying, but its literal translations in, say, French or Mandarin have no such idiomatic meaning. By contrast, a wide variety of languages use the equivalent of *knows with certainty* to denote a state more demanding than run-of-the-mill knowledge. (Thanks to Jonathan Schaffer and Tim Williamson (p.c.) for raising these issues.)

⁷'A case of panic now, not later', http://www.smithers.co.uk/news_article.php?id=16&o=50.

⁸Geeting, *Truckers and Troopers*, p.96

Could one understand epistemic certainty as iterated knowledge, that is, knowing that one knows, or knowing that one knows that one knows?⁹ One difficulty for this suggestion is that epistemic certainty—like subjective certainty—comes in degrees, whereas knowledge does not (Stanley 2005a; Pavese MS):

(46) It's fairly $\left\{ \begin{array}{c} \checkmark \text{ certain} \\ \# \text{ known} \end{array} \right\}$ that the butler did it.

(47) It's 95% $\left\{ \begin{array}{c} \checkmark \text{ certain} \\ \# \text{ known} \end{array} \right\}$ that global warming has human causes.

An adequate account of epistemic certainty will account for the fact that it comes in degrees. However, it's highly doubtful that there's any one-to-one mapping between degrees of certainty and iterations of knowledge. (It's not as though 99% certainty corresponds to knowing that you know, and 99.9% certainty corresponds to knowing that you know that you know, etc.)¹⁰

How should we understand epistemic certainty, if not in terms of knowledge? Let's start by considering how to understand degrees of epistemic certainty. For the purposes of this chapter, I propose to officially take degrees of epistemic certainty as primitive. However, we can get an intuitive grip on the notion by relating it to epistemic justification or evidential support: the more epistemically certain a proposition is for someone, the more justification they have for believing it. (The reason I do not present this as an analysis of degrees of certainty is that I'll ultimately be arguing that degrees of justification and evidential support can be profitably understood in terms of degrees of certainty.)¹¹

⁹See Turri (2010) for the view that, in most contexts, a certainty ascription is roughly equivalent to a second order knowledge ascription. See Avicenna for the view that certainty is to 'know that you know, and to know that you know that you know, *ad infinitum*.' (Quoted in Black 2006.)

¹⁰Another way to attack analyses of certainty in terms of iterated knowledge is via counterexample. Suppose that Mary and Harry both know that Nairobi is in Kenya, but for some reason fail to know that they know it. (Perhaps after the first day of Epistemology 101 they have come to doubt whether knowledge is possible, hence they no longer believe that they know Nairobi is in Kenya.) Still, it seems Mary could know that Nairobi is in Kenya with greater epistemic certainty than Harry if she has stronger grounds for her belief (e.g. Harry's belief is based on testimony, whereas Mary lives Nairobi).

¹¹While I am taking degrees of epistemic certainty as a primitive, I want to leave open the possibility of giving some further, reductive account of degrees of certainty. If the view developed in Chapter Two is on the right track, then the prospects for a reductive first-order account of degrees of epistemic certainty is bleak; however, we could still give a metaepistemological reduction of 'degrees of certainty' talk. To apply the attitudinal framework developed there, ascriptions of degrees of epistemic certainty would be analyzed in terms of degrees of epistemic approval.

We can go on to use our understanding of degrees of epistemic certainty to get a purchase on *knowing with certainty*. To know p with epistemic (as opposed to merely subjective) certainty is to know p with a particularly high degree of epistemic certainty—higher than that typically required for knowledge. Such an account seems plausible in light of the minimal pairs we’ve considered ((42a) vs. (42b); (43a) vs. (43b)).

This account also fits with the traditional understanding of certainty as a particularly exalted form of knowledge. Of course, as noted in §3.1, many writers within this tradition go further, regarding certainty as the epistemic ideal. According to this view, certainty is not merely knowledge that reaches a particularly high threshold of justification or support; it is knowledge that reaches the *highest possible* degree thereof. And this in turn invites the worry that hardly any of our knowledge will attain this lofty status.

In the rest of this section, I examine what I take to be the most sophisticated version of this worry. I argue that this worry turns out to be unwarranted. According to the account I suggest, a belief qualifies as epistemically certain as long as it comes *close enough* to the epistemic ideal. If we allow for contextual variation in what counts as ‘close enough’ to the ideal, we can hold that much of our ordinary knowledge qualifies—in appropriate contexts—as certain.

3.2.2 In Favor of Scarce

Call the thesis that either subjective or epistemic certainty is an extreme rarity, ‘Scarce.’ In my eyes, the most compelling argument for Scarce starts from the observation that *certain*—on both its subjective and epistemic uses—is a maximum standard absolute gradable adjective.

First, some background. Gradable adjectives denote properties that come in degrees. As noted in Chapter Two, they admit of degree modifiers (e.g., *very*, *extremely*) and comparatives (e.g., *x is taller than y*, *p is more certain than q*). Within the class of gradable adjectives, maximum standard absolute gradable adjectives (*max adjectives*, hereafter) form a semantically unified subclass.¹² Typically, max adjectives are characterized as adjectives that require their arguments to have the maximal degree of the denoted property. For

¹²See Unger 1975; Kennedy and McNally 2005; Kennedy 2007.

example, *straight* is a max adjective: something is straight iff it has the maximal degree of straightness. Hence the infelicity of (48):

(48) ? The stick is straight, but it's slightly crooked.

As a number of authors have noted, *certain* functions as a max adjective.¹³ This explains the infelicity of (49) and (50):

(49) ? Fred is certain it's going to rain, but he's not entirely certain it's going to rain.

(50) ? It's certain to rain, but it's not entirely certain to rain.

The interactions between *certain* and degree modifiers provide further evidence that *certain* is a max adjective. Max adjectives tolerate maximizing modifiers (e.g., *completely*, *perfectly*) to a much greater degree than relative gradable adjectives (e.g., *long*, *tall*) (Rotstein and Winter 2004; Kennedy 2007):

(51) The road is completely/perfectly $\left\{ \begin{array}{l} ?? \text{ long} \\ \checkmark \text{ straight} \end{array} \right\}$.

This diagnostic gives further evidence that *certain* is a max adjective:

(52) ✓ Sandy's completely/perfectly certain that the Mets will win.

(53) ✓ It's completely/perfectly certain that the Mets will win.

Once we accept that *certain* is a max adjective, it seems we have an easy argument for Scarce. If *certain* is a max adjective, then whenever S is subjectively certain that *p*, S has the maximal degree of subjective certainty that *p*. Similarly, whenever *p* is epistemically certain (for S), *p* has the maximal degree of epistemic certainty (for S). But this is a high bar—one that precious little of our everyday knowledge seems to meet. Take, for instance, my knowledge that Marseilles is in France. While I'm very confident of this fact, it seems wrong to say that I assign the maximal degree of subjective certainty to this fact. Surely I'm not as certain of this fact as I am of the disjunction: *Either Marseilles is in France or it isn't*. The same goes for epistemic certainty: the disjunction is more certain than either disjunct.

¹³Unger 1975; Lassiter 2010, 2011; Klecha 2012.

3.2.3 Against Scarce

But perhaps we shouldn't be so quick. After all, if asked, I'd readily assert both:

(54) I'm certain that Marseilles is in France.

(55) It's certain that Marseilles is in France.

More generally, people are fairly liberal in their certainty ascriptions: they don't reserve *certain* for a tiny sliver of their knowledge.

Thus, our everyday ascriptions of certainty count against Scarce. Of course, those impressed with the argument in §3.2.2 might simply say that all such ascriptions are false.¹⁴ But this seems like a rather desperate and undesirable maneuver. *Ceteris paribus*, it would be nice if we could find a way to make sentences such as (54) and (55) come out true.¹⁵

3.2.4 The Best of Both Worlds: Max Adjectives Without Scarce

A natural way of reconciling the thesis that *certain* is a max adjective with the truth of (54) and (55) is to relativize the semantic values of gradable expressions to contextually determined standards of precision or levels of granularity.¹⁶ Consider *straight*. Some contexts involve strict standards: when constructing a satellite, a slight dent might preclude an antenna from falling under the extension of *straight*. In more permissive contexts, we can truthfully apply *straight* to the same antenna, since it comes close enough to having the maximum degree of straightness.

Here's one way of developing this idea. Assume a scalar semantics for gradable adjectives, according to which every gradable adjective A is associated with a scale Θ_A , where a scale is a set of degrees totally ordered along some dimension. For any max adjective A_{max} :

Close Enough: 'x is A_{max} ' is true in a context c iff x's degree on Θ_A is close enough for the purposes of c to the maximal element of Θ_A .¹⁷

¹⁴See Unger 1975; Lasersohn 1999.

¹⁵As Unger (1975) observes, the argument for Scarce generalizes far beyond the case of *certain*; it seems to show that almost all of our ordinary utterances involving absolute gradable adjectives are false. Even those who aren't particularly worried about ensuring the truth of (54) and (55) might well recoil from the conclusion that such a large swath of ordinary discourse is false.

¹⁶Lewis (1979b) sketches a response to Unger along these lines. For recent work on levels of granularity, see Krifka 2007; Sauerland and Stateva 2007; van Rooij 2011.

¹⁷Here's one way to formalize this. Let \mathbf{A} be a function that takes an entity and delivers its degree on

A gloss on ‘close enough’: a degree is close enough to the maximal element of the scale, given a context, as long as any difference between the two is *ignorable* (i.e. can be appropriately ignored). In many a conversation, the difference between the antenna’s degree of straightness and the maximum degree of straightness is ignorable.

How can we extend this treatment to *certain*? Assume that subjective uses of *certain* are associated with a scale of degrees from 0 to 1 (the subjective certainty scale). Next, we postulate a function that, given a subject S , a time t , and a proposition p , delivers the degree to which S is subjectively certain that p at t . I’ll call this the ‘subjective certainty function’ (SC), but, as far as I’m concerned, the difference between this and the more familiar notion of a credence function is merely terminological. Similarly, assume that epistemic uses of *certain* are associated with a scale of degrees from 0 to 1 (the epistemic certainty scale), and an epistemic certainty function (EC). Given any S , t , p , EC delivers the degree to which p is epistemically certain for S at t . We then say that a *positive form*¹⁸ subjective certainty ascription, ‘ S is (subjectively) certain that p ’ is true in a context c iff S ’s degree of subjective certainty in p is close enough for the purposes of c to 1. Likewise, a positive form epistemic certainty ascription, ‘ p is (epistemically) certain (for S)’ is true in c iff the degree to which p is epistemically certain for S is close enough for the purposes of c to 1.¹⁹

This enables us to reject Scarce: in most contexts, the degree of subjective and epistemic certainty of *Marseilles is in France* is close enough to the maximum. Hence most utterances of (54) and (55) are true. However, we can still explain the data that motivated the classification of *certain* as a max adjective. Take (50) (repeated here as (56)):

Θ_A . Let **max** be a function that takes a scale and delivers the highest degree on that scale. Finally, let g be a function that, for any entity x , context c , scale Θ_A , and degree $d \in \Theta_A$ gives you the set of all degrees Γ s.t. if $\mathbf{A}(x) \in \Gamma$, this counts as close enough to d for the purposes of c . We then give the following lexical entry for a max adjective A_{max} :

$$[\text{is } A_{max}]^c = [\lambda x_e . \mathbf{A}(x) \in g(x, c, \Theta_A, \mathbf{max}(\Theta_A))]$$

Note that this way of formulating the semantics allows that what’s ‘close enough’ can vary with x : within a given context, the standards for a full wine glass might differ from the standards for a full gas tank. (See Foppolo and Panzeri 2010 for experimental data indicating that judgments about the applicability of a max adjective to an object are sensitive to the object in question.)

¹⁸The ‘positive’ or unmarked form of a gradable adjective lacks any overt degree morphology. Compare, for example, *I’m certain* (positive) with *I’m fairly/entirely/95% certain* (marked).

¹⁹One way to formalize this would be to offer separate lexical entries for the two uses of *certain*:

$$\begin{aligned} [\text{is } \text{certain}_{subj}]^c &= [\lambda p_{\langle s, t \rangle} [\lambda t_i [\lambda x_e . \text{SC}_{x, t}(p) \in g(\langle x, t, p \rangle, c, \Theta_{SC}, \mathbf{max}(\Theta_{SC}))]]] \\ [\text{is } \text{certain}_{epi}]^c &= [\lambda p_{\langle s, t \rangle} [\lambda t_i [\lambda x_e . \text{EC}_{x, t}(p) \in g(\langle x, t, p \rangle, c, \Theta_{EC}, \mathbf{max}(\Theta_{EC}))]]] \end{aligned}$$

(Here Θ_{SC} is the subjective certainty scale and Θ_{EC} is the epistemic certainty scale.)

(56) ? It's certain to rain, but it's not entirely certain to rain.

Assume that the maximizing modifier *entirely* is also sensitive to standards: feed it an adjective A and it will deliver the set of degrees on Θ_A that are close enough to the maximum. Then (56) is inconsistent, provided we hold the standards fixed: the first conjunct asserts that *It will rain* comes close enough to the maximal degree of epistemic certainty; the second conjunct denies this. At the same time, our account has the resources to explain why (56) doesn't sound *as* bad as overt contradictions such as:

(57) # It's certain to rain, but it's not certain to rain.

After all, one effect of maximizing modifiers is to raise the standards of precision (Sassoon and Zevakhina 2012). And this makes a 'repair' reading, according to which the second conjunct is interpreted using heightened standards, more readily available for (56) than for (57).

3.2.5 Puzzling Pos

One might worry that our semantics has difficulty explaining some of the entailments of positive form (*pos form*, hereafter) epistemic certainty ascriptions. In particular, one might worry that it has trouble explaining why such ascriptions are factive (Moore 1959):

- (58) a. It's certain that the Mets will win. \Rightarrow
 b. The Mets will win.²⁰

Suppose that Sal has good but not conclusive evidence that the Mets will win: for him, the degree of epistemic certainty that the Mets will win is .8. Furthermore, suppose that Sal is in a conversation governed by low standards of precision (he's in a pub, not an epistemology seminar). Still, if Sal asserts (58a) and the Mets in fact lose, it seems that Sal said something false. To see this, consider how natural it would be for a chagrined Sal

²⁰Like many philosophers, I call a verb 'factive' if it entails the truth of its complement clause. Semanticists would prefer the term 'veridical', reserving the term 'factive' for a verb that presupposes the truth of its complement clause.

to retract his assertion: ‘I guess I was wrong when I said that it was certain that the Mets will win.’²¹

How can our scalar semantics explain the factivity of pos form epistemic certainty ascriptions? One could try building factivity into the epistemic certainty scale: perhaps whenever $EC_{S,t}(p) \geq .8$, p is guaranteed to be true. But the entailment patterns of graded certainty ascriptions cast doubt on this strategy:

- (59) a. It’s 99.9% certain that the Mets will win. \nRightarrow
 b. The Mets will win.

Even if (59a) is uttered in a context with high standards, we can’t infer (59b).

We face a seemingly inconsistent triad: (i) no degree of epistemic certainty less than 1 suffices for factivity; (ii) in some contexts, the degree of certainty required for pos form epistemic certainty ascriptions to be true is less than 1; (iii) in all contexts, pos form epistemic certainty ascriptions are factive. I suggest resolving this apparent inconsistency by imposing constraints on when a degree can be close enough to the maximum. In particular, I suggest a Truth Constraint: if p is false, it’s never appropriate to ignore the difference between p ’s degree of epistemic certainty and the maximum. And so no degree of epistemic certainty shy of 1 counts as close enough to 1. If p is true, greater latitude is permitted.²²

We can handle other puzzling entailments of pos form epistemic certainty ascriptions in similar fashion. Arguably, pos form epistemic certainty ascriptions don’t merely entail the truth of the embedded proposition; they also entail that the embedded proposition is known. The linguistic data here is not entirely clear-cut, and intuitions may vary across speakers. However, it seems to me that in any situation where a knowledge ascription is false, the corresponding pos form epistemic certainty ascription is also false. Consider a standard Gettier case: Jessie comes to believe that it’s 3pm on the basis of a stopped clock. It just so happens that it is 3pm (Russell 1948). If Jessie were to claim that it’s certain

²¹By contrast, pos form subjective certainty ascriptions are not factive: *I’m certain the Mets will win* \nRightarrow *The Mets will win*.

²²Readers may note a resemblance between this proposal and Lewis’ (1996) treatment of *knows*, which captures the entailments of knowledge ascriptions via constraints on when a possibility is properly ignored. In particular, the Truth Constraint resembles Lewis’ ‘Rule of Actuality’, according to which the actual world is never properly ignored.

that it's 3pm, I think there's some tendency to regard her utterance as false. Likewise in lottery cases: suppose we all believe, on statistical grounds, that a given lottery ticket T is a loser. Suppose that T does indeed lose. Still, it seems that there's an inclination to deny that, prior to the draw of the winning ticket, it was certain that T was a loser.

Our strategy for handling factivity extends to handle the knowledge entailment. Let C be any necessary condition on knowledge. If no degree of epistemic certainty shy of 1 guarantees that C obtains, we can preserve our scalar semantics for *certain* by converting C into a constraint on the 'close enough' relation: if C isn't satisfied, no difference between p 's degree of epistemic certainty for S and the max is ignorable. For example, suppose one adopts the popular view that a safety requirement on knowledge explains the intuition that the protagonists of Gettier and lottery cases lack knowledge.²³ We could then convert this requirement into a Safety Constraint: if S 's belief in p could have easily been held falsely, there's no context in which a difference between p 's degree of epistemic certainty for S and the maximum is appropriately ignored.²⁴

To illustrate the strategy, suppose *The Mets will win* is epistemically certain for both Sue and Sal to degree .8. Sue is in the good case: for any necessary condition on knowledge C , Sue's belief that the Mets will win satisfies C . Sal is in the bad case: his belief that the Mets will win doesn't satisfy C . In Sue's context, a degree of .8 epistemic certainty counts as close enough to the maximum; hence she can truly assert (58a). In Sal's context, no degree of epistemic certainty shy of 1 counts as close enough to the max; hence his utterance of the same sentence is false. (See Figure 3.1.)²⁵

²³See Sosa 1999, 2004; Williamson 2000a; Pritchard 2005; Lasonen-Aarnio 2010.

²⁴Of course, the idea that knowledge requires safety is controversial. (For putative counterexamples, see Neta and Rohrbaugh 2004; Comesaña 2005; Kelp 2009; Bogardus 2014. For the view that many of the counterexamples can be avoided by refining our understanding of 'close' possibilities, see Williamson 2009.) However, for present purposes the Safety Constraint merely serves as a useful illustration of a more general strategy: whatever one's preferred conditions on knowledge, we can construe these conditions as constraints on the 'close enough' relation. In doing so, we'll be able to reconcile the knowledge entailment with our scalar semantics for *certain*.

²⁵The strategy of imposing substantive constraints on the 'close enough' relation could also be used to underwrite a multi-premise closure principle on epistemic certainty (e.g., if (i) $p_1 \dots p_n$ are certain, (ii) $p_1 \dots p_n$ entail q , (iii) one competently deduces q from $p_1 \dots p_n$, then q is also certain). To reconcile such a principle with a scalar semantics for *certain*, one can convert one's preferred solution to the lottery paradox into a constraint on the 'close enough' relation. For example, Leitgeb (2014) proposes to solve the lottery paradox via a 'stability' constraint on belief. According to this approach, for S to believe p is for S 's credence in p to remain sufficiently high even when conditionalized on any proposition that is both consistent with p and assigned some non-zero credence by S . Those attracted to Leitgeb's solution could transpose it into our framework: in order for p 's degree of epistemic certainty to be close enough to the max, p 's degree of




Figure 3.1: Good vs. Bad.

Let me briefly address two worries one might have about the current strategy for reconciling our scalar semantics for *certain* with the entailments of pos form epistemic certainty ascriptions. First, one might worry that this strategy presupposes that epistemic certainty is analyzable—that it’s possible to state non-circular necessary and sufficient conditions for the truth of pos form epistemic certainty ascriptions. If so, isn’t this presupposition overly optimistic? After all, recently some philosophers have suggested that the analysis of knowledge is a quixotic endeavor.²⁶ If knowledge is unanalyzable, then, given the knowledge entailment, it’s natural to worry that epistemic certainty is similarly unanalyzable.

However, even if epistemic certainty turned out to be unanalyzable, this wouldn’t undermine the current strategy. What this would mean (on the current approach) is that we won’t be able to give any non-circular statement of the conditions under which some degree of epistemic certainty counts as close enough to the max. While this might be disappointing for those with reductive aspirations, it does not pose any problems for our basic account of certainty. After all, we’ll still be able to maintain Close Enough. In doing so, we’ll still be able to accept that *certain* is a max adjective while resisting the siren song of Scarce.

Second, one might worry that even if we can account for the entailments of pos form epistemic certainty ascriptions by imposing substantive constraints on the ‘close enough’

epistemic certainty must remain sufficiently high even when p is conditionalized on any proposition q such that (i) q is consistent with p , (ii) $EC(q) > 0$.

²⁶See, in particular, Williamson 2000a.

relation, this strategy proves ad hoc. By contrast, an austere semantics for max adjectives that dispenses with standards of precision and embraces Scarce (Unger 1971, 1975; Kennedy 2007) seems like it can avoid such ad hocery. Perhaps this gives us reason to reconsider our rejection of Scarce.

But closer inspection gives reason to doubt that austere views have any advantage here. Even those who accept Scarce will presumably admit that in ordinary contexts speakers are willing to assert (55), and listeners are willing to regard this assertion as true. Even if this is just ‘loose talk’, any adequate account of our linguistic practice will include a theory of this ‘loose talk’, and this theory will include a specification of the conditions under which we’re willing to regard (55) as true. Presumably, any such theory will look a lot like Close Enough, recast as a pragmatic story: we regard an utterance of ‘It is certain that p ’ as true iff p ’s degree of epistemic certainty (for the speaker) is close enough to the max.²⁷ Such theories will still face the difficult task of explaining why, in ordinary contexts, we regard an utterance of ‘It is certain that p ’ as false if p is false (and, more generally, if p isn’t known). It is unclear how they could accomplish this task without imposing substantive constraints on the ‘close enough’ relation.

3.2.6 Taking Stock

I’ve argued that we should resist two impulses: the impulse to analyze certainty in terms of knowledge, and the impulse to dismiss certainty as unattainable. According to the treatment of certainty offered here, certainty comes in two forms: subjective and epistemic. The former consists in a strong conviction; the latter consists in a strong epistemic position, not reducible to knowledge. What’s more, while *certain* is a max adjective, this does not entail Scarce: in many contexts, a non-negligible subset of our everyday knowledge qualifies as both subjectively and epistemically certain.

By showing that certainty is not scarce, we’ve paved the way for putting certainty to work in epistemological theorizing. But before getting down to brass tacks, it’s worth highlighting that my account of certainty really involves two sets of distinctions. In addition

²⁷The theory of ‘pragmatic halos’ in Lasersohn (1999) can be understood as a pragmatic story along these lines.

to the subjective/epistemic certainty distinction, there's the distinction between *degrees of certainty* and certainty *full-stop* (the sort that's picked out by pos form ascriptions). In what follows, these different facets of certainty will serve different explanatory roles. Degrees of epistemic certainty are used to explain evidential probability (§3.3), whereas certainty full-stop is used to explain epistemic modals (§3.4) and knowledge (§3.5). Finally, both degrees of certainty and certainty full stop feature in the account of the normative constraints on action (§3.6).

Does this mean the Certainty First program is fundamentally disjunctive? I think not. Degrees of certainty and certainty full-stop are two sides of the same coin. If the discussion in §§3.2.4-3.2.5 is correct, certainty full-stop can be understood in terms of degrees of certainty that come close enough to the maximum. And even if this analysis requires revision, it seems extremely plausible that our grasp on certainty full-stop depends on our awareness that certainty comes in degrees. The fact that we'll be exploiting the full explanatory resources that certainty offers should not obscure the underlying unity of the resulting account.

3.3 Evidential Probability

In everyday conversation, we frequently use probability operators: expressions such as *likely* and *probable*. One common use of such operators is to convey *evidential probability*: the probability of a hypothesis given a body of evidence. For example:

(60) It's 99% likely that the Mets will win.

(60) has a natural reading on which it says that, given the available evidence, there's a 99% chance that the Mets will win.

In this section, I argue that evidential probability ascriptions are best analyzed in terms of epistemic certainty. I begin by providing linguistic data that establish a close connection between evidential probability ascriptions and epistemic certainty ascriptions. A natural explanation for this connection is that evidential probabilities are degrees of certainty (§3.3.1). Not only does this hypothesis explain our linguistic practice, it also does epistemological work: I show that the resulting account of evidential probability has important epistemolog-

ical advantages over rival accounts that explain evidential probability in terms of knowledge (§3.3.2).

3.3.1 Evidential Probabilities and Degrees of Certainty

Both *likely* and *probable* accept proportional modifiers (e.g., *99%*, *95%*). And when both are embedded under the same proportional modifier, they seem to be equivalent. To see this, compare (60) with (61):

(61) It's 99% certain that the Mets will win.

(60) and (61) seem interchangeable, at least when (60) is interpreted in terms of evidential (as opposed to, say, objective) probability. Indeed, it would be quite odd to affirm one while denying the other:

(62) ?? It's 99% certain that the Mets will win. But it's only 98% likely that they'll win.

(63) ?? It's 98% certain that the Mets will win. But it's 99% likely that they'll win.

To my ears, (62) and (63) are only coherent if we impose some non-evidential interpretation on the probability operators.

This close connection cries out for explanation. A natural explanation is that *evidential probabilities simply are degrees of epistemic certainty*. Here's one way of spelling this out. Following recent work on the semantics of probability operators, let us hypothesize that the meaning of a probability operator can be modeled by a probability function.²⁸ Let $\text{Pr}_{S,t}$ be a probability function that, for any proposition p , assigns p a number in $[0, 1]$ representing that probability of p given S 's total evidence at t .²⁹ I propose that we can simply identify Pr with the epistemic certainty function EC :

Evidential Probabilities are Degrees of Epistemic Certainty (EP=EC):

For any proposition p , any subject S , and any time t , $\text{Pr}_{S,t}(p) = \text{EC}_{S,t}(p)$.

²⁸See Swanson 2006; Yalcin 2007, 2010, 2011; Lassiter 2010; Lassiter and Goodman 2015; Moss 2015b.

²⁹Relativization to subjects is needed, since something that is likely given one person's body of evidence can be unlikely given another's. Where it won't cause confusion, I will omit the subscripts.

EP=EC explains why (60) and (61) are interchangeable, and why (62) and (63) are incoherent. This strikes me as an important mark in its favor. What's more, as we'll see shortly, EP=EC enjoys epistemological advantages over rival accounts of evidential probability in terms of knowledge. But before turning to a comparison with the Knowledge First account, I want to tackle a natural objection that may seem, at first blush, to render hopeless any reduction of evidential probabilities to degrees of certainty.

According to the objection in question, I've cherry-picked my data: the equivalence between graded epistemic certainty ascriptions and graded evidential probability ascriptions only holds for proportional modifiers that denote very high degrees on the corresponding scale. But when we look at mid-scale proportional modifiers (e.g., *60%*) the equivalence breaks down:

$$(64) \quad \text{It's } 60\% \left\{ \begin{array}{l} \checkmark \text{ likely} \\ ? \text{ certain} \end{array} \right\} \text{ the Mets will win.}$$

While it's perfectly felicitous to talk of 60% probability, talk of 60% certainty sounds bizarre.

The objection can be generalized by looking at other modifiers. We've seen that *certain* accepts maximizing modifiers. However, probability operators do not:

$$(65) \quad ? \text{ It's completely likely/probable the Mets will win.}$$

Should we conclude that EP=EC must be abandoned? I think not. We've already seen that a max adjective targets the upper end of its scale. What (71) reveals is that our tendency to reserve *certain* for the upper end of the epistemic certainty scale persists even under proportional modifiers: we are happier with *99% certain* than *70% certain*, and less happy still with *60% certain*.

By contrast, *probable* and *likely* are relative gradable adjectives: they do not target the upper end of the probability scale, but rather some contextually determined point along it.³⁰ Thus if there's a 70% chance that it will rain today and an 80% chance that it will rain tomorrow, one can say:

$$(66) \quad \text{It's likely to rain today, but it's more likely that it will rain tomorrow.}$$

³⁰Cf. Yalcin 2010: p.930.

This explains why *probable* and *likely* do not accept maximizing modifiers: as we saw in §3.2.2, relative adjectives resist maximizing modifiers. And since *probable* and *likely* have no tendency to target the upper end of their scale, they happily combine with mid-scale proportional modifiers.

The suggestion, then, is that two adjectives can denote points on the same scale, even if those two adjectives tolerate different modifiers. In case this seems *ad hoc*, it's worth noting that there is independent reason to accept this suggestion. Presumably, *filthy* and *dirty* denote points on the same scale: the difference between being dirty and being filthy is a difference in degree, not in kind. But they accept different modifiers:

$$(67) \text{ Your shirt is completely } \left\{ \begin{array}{c} \checkmark \text{ filthy} \\ ? \text{ dirty} \end{array} \right\}.$$

$$(68) \text{ Your shirt is slightly } \left\{ \begin{array}{c} ? \text{ filthy} \\ \checkmark \text{ dirty} \end{array} \right\}.$$

We can also illustrate this point using ‘extreme adjectives’ (e.g., *gigantic*, *huge*). Intuitively, *gigantic* occupies the same scale as *big*: to be *gigantic* is just to be extremely *big*. But *gigantic* doesn't accept the same modifiers as *big*:

$$(69) \text{ King Kong is very/extremely } \left\{ \begin{array}{c} \checkmark \text{ big} \\ ?? \text{ gigantic} \end{array} \right\}.^{31}$$

Thus we can maintain EP=EC by denying the underlying assumption that two adjectives can only be scale-mates if they accept the same modifiers—an assumption that we have independent reason to reject ((67)-(69)).³² On the picture that emerges, absolute and relative gradable adjectives can co-habit the same scale; when they do, we should expect them to tolerate different modifiers.

Adopting this picture enables us to explain further inferences linking evidential probability talk and epistemic certainty ascriptions. For example, pos form evidential probability ascriptions typically implicate that the corresponding pos form epistemic certainty ascription doesn't hold:

³¹See Morzycki (2012) for a degree-based semantics for extreme adjectives designed to explain this difference in modifier distribution.

³²For further arguments against the view that scale-mates always accept the same modifiers, see Lassiter (2010); Lassiter and Goodman (2015).

- (70) a. It's likely the butler did it \rightsquigarrow
 b. It's not certain the butler did it.

As Lassiter (2010) observes, this seems to be a scalar implicature, akin to the implicature from *some* to *not all*. And this fact is naturally explained by holding that while *certain* is a max adjective and *likely* is a relative adjective, they nonetheless occupy the same scale.

3.3.2 Epistemological Advantages

One task for a theory of evidential probability is to capture the meaning of everyday evidential uses of probability operators. Another task is to provide an *epistemologically serviceable* account of evidential probability—that is, an account that can be put to use in epistemological theorizing. It is not *a priori* obvious that these two tasks will coincide: it may turn out that our everyday notion of evidential probability is ill-suited to perform the jobs that epistemologists demand of it.

Thus far, I've focused on the first task, arguing that an analysis of evidential probability in terms of degrees of epistemic certainty sheds light on the close connections between evidential probability talk and epistemic certainty talk. I now turn to the second. I show that by understanding evidential probabilities in terms of degrees of certainty, we avoid two serious obstacles for 'Knowledge First' accounts of evidential probability.

In order to lay out the advantages, I begin by reviewing the Knowledge First treatment of evidential probability. Williamson famously proposes that evidence is knowledge:

E=K: For any subject *S* and any time *t*, *S*'s total evidence at *t* = {*p*: *S* knows *p* at *t*}.³³

Williamson goes on to suggest that the evidential probability of a proposition *p* is the probability of *p* conditional on one's total evidence (2000a: chp.10). Let $K_{S,t}$ be the conjunction of everything *S* knows at *t*. Williamson's account amounts to the following:

Knowledge Account of Evidential Probability:

$$\Pr_{S,t}(p) = \Pr_{S,t}(p|K_{S,t}), \quad \text{where } \Pr_{S,t}(p|K_{S,t}) \neq 0.^{34}$$

³³See Williamson 1997, 2000a: chp.9.

³⁴This is Williamson's synchronic characterization of evidential probability. He also offers a diachronic

An immediate consequence of this is that all knowledge has evidential probability 1:

Knowledge has Evidential Probability 1: $\Pr_{S,t}(p) = 1$ if S knows p at t .

While the Knowledge Account is deservedly influential, this particular consequence gives rise to two serious difficulties—difficulties that a Certainty First account of evidential probability neatly avoids. I will lay out each difficulty in turn.

At least one goal of an epistemologically serviceable theory of evidential probability is to model our intuitive notion of *strength of epistemic position*: the stronger my epistemic position with regards to p , the higher its evidential probability (for me) should be. However, it seems that I stand in a better epistemic position with regards to some propositions I know than I do towards others. Consider the following propositions:

Snow: It will snow in NY sometime in the next year.

Snow or no snow: Either it will snow in NY sometime in the next year or it won't.

Assuming inductive skepticism is false, I know both propositions. But, intuitively, I stand in a stronger epistemic position towards *Snow or no snow* than I do towards *Snow*. And so *Snow or no snow* should receive a higher evidential probability than *Snow*. But this conflicts with the Knowledge Account, which assigns both propositions probability 1.

If we take evidential probabilities to be degrees of epistemic certainty, we avoid this difficulty. After all, I can know some propositions with a higher degree of epistemic certainty than others. In particular, *Snow or no snow* has a higher degree of epistemic certainty for me than *Snow* does. Given $EP=EC$, $\Pr(\textit{Snow or no Snow}) > \Pr(\textit{Snow})$.³⁵

The second difficulty for the Knowledge Account comes from considering the relation between evidential probabilities and credences. At least as far back as Locke's *Essay*, epistemologists have been attracted to the view that rationality requires one proportion

characterization, according to which there is an initial evidential probability function that reflects the 'intrinsic plausibility of various hypotheses' (Williamson 2000a: 211), which is then updated by conditionalization as new evidence is acquired. Here I focus on the synchronic aspect of his account.

³⁵It is worth highlighting a closely related advantage of the Certainty First account. Many epistemologists have been attracted to the fallibilist idea that one can know p even if the evidential probability of p is less than 1. The Knowledge Account precludes fallibilism, thus construed. By contrast, $EP=EC$ makes room for fallibilism: fallible knowledge will be present whenever p 's degree of epistemic certainty is high enough to allow for knowledge, while still falling short of the maximum.

one's degree of belief to the evidence.³⁶ This suggests that an epistemologically serviceable account of evidential probability will provide normative constraints on credences. Here's a natural way of formulating such a constraint more precisely:

Credal Constraint: For any S, t, p : S 's credence in p at t should equal $\Pr_{S,t}(p)$.

Why does this cause trouble for the Knowledge Account? Intuitively, I should be more confident of some of the things I know than others. For example, my credence in *Snow or no snow* should be 1, whereas it seems that my credence in *Snow* should be slightly lower. Given the Credal Constraint, it follows that the evidential probability of the disjunction should be higher than the probability of the disjunct. But once again this is inconsistent with the Knowledge Account, which assigns both propositions probability 1.³⁷

Here too, taking evidential probabilities to be degrees of certainty enables us to avoid the difficulty. We've seen that $EP=EC$ allows for $\Pr(\textit{Snow or no snow})$ to exceed $\Pr(\textit{Snow})$, since the former is more epistemically certain than the latter. From the Credal Constraint, it follows that my credence in the former should exceed my credence in the latter.

Indeed, combining the Credal Constraint with $EP=EC$ delivers an elegant picture of the normative links between the two species of certainty, epistemic and subjective. As noted in §3.2.4, the subjective certainty function seems to be equivalent to a credence function. If this is right, conjoining $EP=EC$ with the Credal Constraint entails:

Matching Requirement: Your degree of subjective certainty in p should equal the degree to which p is epistemically certain for you.

We can provide independent motivation for the Matching Requirement. Suppose that it's 99% certain that the Mets will win. What should your credence be that the Mets will win? There seems to be an obvious answer: .99. Consider how odd the following sounds:

(71) ?? It's 99% certain that the Mets will win. But I'm only 98% certain that they'll win.

³⁶Williamson also endorses this requirement (2000a: p.223).

³⁷See Kaplan 2003, 2009 for a closely related objection.

The Matching Requirement explains these intuitions: (71) is odd because no one could truly assert it without violating a basic requirement of rationality.

3.3.3 Further Issues

I've laid out the bare bones of a theory of evidential probability. On my theory, certainty plays a starring role: evidential probabilities are degrees of epistemic certainty, and these impose normative constraints on our degrees of subjective certainty. In the course of developing my theory, I skirted a number of points of detail. Here I'll briefly take up two such points.

First, I haven't yet taken a stand on what's involved in *possessing evidence*. There are various ways of developing a theory of evidence possession within the present framework, but one attractive option is to adopt a *gradational* conception. According to a gradational conception, whether p is part of a subject's evidence isn't an all or nothing matter; instead, it comes in degrees. Joyce articulates the idea thus:

On a categorical conception, the question of whether a belief has evidential standing has an unqualified yes/no answer... On a gradational view, one speaks not of evidential status *tout court*, but of the degree to which a believed proposition counts as evidence. Evidential status falls along a spectrum that ranges from the best sort of evidence, through intermediate grades, to beliefs that are not evidence at all. (Joyce 2004: 298-299)

As Williamson (2004) observes, gradationalists owe us a story about how to understand differences in evidential status. The present perspective provides such a story: degrees of evidential status just are degrees of epistemic certainty. Thus, if p is evidence for S at t to degree .6, and q is evidence for S at t to degree .7, this is because $EC_{S,t}(p)=.6$ and $EC_{S,t}(q)=.7$.³⁸

A second point of detail concerns the diachronic dimension of evidential probability: are there any update rules governing how evidential probabilities evolve over time? This

³⁸Arguably, this conception of degrees of evidence was anticipated by Hume: Schmitt interprets Hume as identifying degrees of evidence with "grades of certainty" (2014: 77).

is a vexed question, and the present account is compatible with either a negative or an affirmative answer. If we answer in the negative, the result is a ‘time-slice epistemology’ (Hedden 2015; Moss 2015a), according to which the only rational constraints on credences are synchronic (specifically, the constraints provided by the Matching Requirement). If we answer in the affirmative, then we can avail ourselves of any number of different update rules in the literature.

To give just one example, the present account could be combined with the view that the evidential probability function is updated via Jeffrey Conditionalization (Jeffrey 1965). To see how this would work, let us define *S*’s *unimpeachable evidence* at *t* as the set of propositions that are maximally certain for *S* at *t* ($\{p : EC_{S,t}(p) = 1\}$). Let E_i be a partition over the worlds compatible with *S*’s unimpeachable evidence at *t*. Jeffrey Conditionalization amounts to the following update rule:

$$\textbf{Jeffrey Conditionalization: } \Pr_{S,t_2}(p) = \sum_i \Pr_{S,t_1}(p \mid E_i) \Pr_{S,t_2}(E_i)$$

While the account of evidential probability defended here is by no means committed to Jeffrey Conditionalization, it does provide Jeffrey Conditionalizers with resources to defuse at least one persistent criticism—specifically, that we lack an independent account of the probabilities in question. As Williamson (2000a) puts it: “Jeffrey conditionalization is hard to integrate with any adequate epistemology, because we have no substantive answer to the question: what should the new weights [i.e. $\Pr_{(S,t_2)}(E_i)$] be?” (2000a: 216) The present approach gives a substantive answer to this question: it says that the new weights are degrees of epistemic certainty.³⁹

3.3.4 Objections

I’ve argued that evidential probabilities should be understood in terms of degrees of certainty. The resulting analysis explains the close connections between epistemic certainty

³⁹Of course, Jeffrey Conditionalization faces further challenges. Perhaps the most well-known worry is that Jeffrey Conditionalization is non-commutative: the order in which one receives new evidence affects how evidential probabilities get updated. (For versions of this objection, see Levi 1967; Domotor 1980; Skyrms 1986; van Fraassen 1989. For an argument that the non-commutativity countenanced by Jeffrey Conditionalization is unproblematic, see Lange 2000.) Since my account is not committed to Jeffrey Conditionalization, I will avoid taking a stand on whether such further challenges can be resolved.

ascriptions and evidential uses of probability operators (§3.3.1); it also enjoys important advantages over the Knowledge Account (§3.3.2). In the rest of this section, I defend my account against two natural objections.

First, one might question my account's supposed advantages over the Knowledge Account. Didn't I say that my knowledge that *Marseilles is in France* counts as epistemically certain, at least in ordinary contexts? And if both *Marseilles is in France* and *Marseilles is in France or it isn't* qualify as epistemically certain in a context, how do we avoid assigning them the same evidential probability?

In response, it's important to emphasize that, according to my view, a proposition p 's evidential probability does not depend on whether p falls under the extension of *certain* in a context; rather, it depends on p 's position on the epistemic certainty scale. Thus while *certain* is context-sensitive, evidential probabilities are not.⁴⁰ Even though in most contexts *Marseilles is in France* will qualify as certain, its degree of certainty will still be less than 1. And so my account can capture the fact that $\Pr(\textit{Marseilles is in France}) < \Pr(\textit{Marseilles is in France or it isn't})$.

One might worry that the objection is not so easily evaded. After all, my account makes use of an evidential probability function, which presumably obeys the Kolmogorov axioms. So won't all logical truths be assigned evidential probability 1? If so, then it seems the problem re-emerges: intuitively, my epistemic position with respect to simple tautologies is stronger than my epistemic position with respect to complex logical theorems. Similarly, it seems I should have a higher credence in the former than the latter.

While I agree that there's at least a *prima facie* difficulty here, I think it's important to be clear about the source of the problem. The problem doesn't stem from my claim that our intuitive notion of evidential probability maps onto the notion of degrees of certainty. Even when it comes to logical truths, these two notions march hand-in-hand: intuitively, a complex logical conjecture that turns out to be true has a lower evidential probability than a simple tautology. ('This conjecture is likely to be true, but we won't know for sure

⁴⁰Of course, if we adopt a contextualist attitudinal semantics for "degrees of certainty" talk, then this isn't quite right, since different speakers will assert different things when they say things like: "It's .99% certain for S that p ." Still, while there will be *inter*-subject variations in the extensions of evidential probability ascriptions, there will not be *intra*-subject variations.

until we have a proof.’) Rather, the problem arises from thinking that our ordinary notion of evidential probability can be modeled using a probability function that treats all logical truths on a par.

This problem is closely connected to the familiar problem of logical omniscience, and the literature on the problem of logical omniscience offers various strategies that—if successful—could be borrowed to resolve our present difficulty. For instance, one could borrow Garber (1983)’s strategy and assign evidential probabilities not to subsets of the logically possible worlds, but rather to sentences in a language generated from the logical closure of certain atomic sentences. Or one could take a page from Stalnaker (1999): for every logical truth L that we intuitively want to assign an evidential probability less than 1, we could take our intuitions to be tracking the evidential probability of the diagonal of some sentence that actually expresses L .⁴¹ But note that neither of these strategies will help rescue the Knowledge Account of Evidential Probability, since it is a direct consequence of the Knowledge Account that all knowledge gets assigned evidential probability 1.

A second objection to my account focuses on my proposed direction of explanation. Suppose we grant that the evidential probability of a proposition equals its degree of epistemic certainty. Still, the objection runs, it doesn’t follow that degrees of certainty are explanatorily prior to evidential probabilities. Perhaps evidential probabilities should be used to explain degrees of certainty, or perhaps neither has priority.

Questions about the direction of explanation are difficult to adjudicate. In my eyes, the main reason for being attracted to the direction of explanation advocated here is that epistemic certainty can be used to illuminate a wide variety of epistemic phenomena, not just evidential probability.⁴² Of course, it’s one thing to assert that certainty can be used to explain various epistemic phenomena, it’s another thing to show it. It’s to this task that I now turn. In the rest of this chapter, I use certainty to shed light on epistemic modals, knowledge, and the normative constraints on action. In doing so, I aim to build a strong

⁴¹The diagonal of a sentence ϕ is the proposition that, for any world w in the context set, is true at w iff the proposition expressed by ϕ at w is true at w (i.e., $\lambda c . \llbracket \phi \rrbracket^{c, w_c} = 1$).

⁴²Compare with the Knowledge Account: while Williamson does not offer E=K as a conceptual analysis, he does take knowledge to explain evidence, rather than the other way around (2000a: 185-186). And presumably one reason for taking this to be the correct direction of explanation is that (according to Knowledge Firsters) we can use knowledge to explain so much else in the epistemic domain.

case that certainty plays a central explanatory role in epistemology more generally.⁴³

3.4 Epistemic Modals

3.4.1 Two Analyses of Epistemic Modals

We often use modals (*might*, *possibly*, *must*, etc.) to convey information about our epistemic state. Sometimes these epistemic modals occur with restrictor phrases (e.g., *In light of the evidence, he must be the culprit*), but they frequently occur *bare*, as in:

(72) It might/must be raining.

The standard analysis of bare epistemic modals treats them as quantifiers over the possibilities compatible some relevant epistemic state:

‘Might ϕ ’ is true at a point of evaluation i iff ϕ is compatible with the relevant epistemic state.

‘Must ϕ ’ is true at i iff the relevant epistemic state entails ϕ .⁴⁴

Much of the literature on epistemic modals focuses on *whose* epistemic state is relevant. Is it some group determined by the context of utterance, as contextualists maintain? Or is the relevant group determined, at least in part, by the context of assessment, as relativists insist? For present purposes, I want to set this question aside, and focus on the question of what *sort* of epistemic state is relevant.

The most common view is that the relevant epistemic state is knowledge. Call this the ‘Knowledge Analysis’:

‘Might ϕ ’ is true at i iff ϕ is compatible with what the relevant agents know (or are in a position to know).

⁴³Note that even those who are unconvinced by my claims about the direction of explanatory priority can still accept much of what is said in this section. Even if we deny that either side of EP=EC is explanatorily prior to the other, EP=EC still provides an important and interesting connection between certainty and evidential probability. By establishing this connection, we’ll have already provided some motivation for the thesis that certainty plays a relatively central role in the epistemic realm.

⁴⁴For the sake of simplicity, I omit the role of the ordering source. I’ll also ignore the fact that some take ‘might’ and ‘must’ to carry presuppositions of indirectness—see e.g. von Fintel and Gillies 2010.

‘Must ϕ ’ is true at i iff ϕ is entailed by what the relevant agents know (or are in a position to know).⁴⁵

By contrast, I advocate a ‘Certainty Analysis’, according to which the relevant state is certainty.⁴⁶ There are a couple of ways of developing a Certainty Analysis, depending on whether one thinks that epistemic modals are sensitive to agents’ states of epistemic certainty, their states of subjective certainty, or both. For present purposes, I’ll operate with a simple account, according to which epistemic modals are analyzed only in terms of agents’ states of epistemic certainty:

‘Might ϕ ’ is true at i iff ϕ is compatible with the set of propositions that are epistemically certain for the relevant agents.

‘Must ϕ ’ is true at i iff ϕ is entailed by the set of propositions that are epistemically certain for the relevant agents.⁴⁷

In what follows, I motivate the Certainty Analysis by pointing to linguistic data that suggest a close connection between epistemic certainty and epistemic modals. The Certainty Analysis illuminates this connection, whereas the Knowledge Analysis leaves it unexplained.

3.4.2 In Favor of the Certainty Analysis

Suppose a detective asserts:

⁴⁵For endorsements of the Knowledge Analysis of Epistemic Modals, see Hacking 1967; Kratzer 1981, 2012; DeRose 1991; Egan et al. 2005; Stanley 2005b; Stephenson 2007a; Hawthorne 2007, 2012; Egan and Weatherson 2011; Dorr and Hawthorne 2013.

⁴⁶Only a couple of authors have entertained something like the Certainty Analysis. DeRose suggests that *might* is the dual of *it is certain that* (1998; 2009: 20). However, he seems to think certainty should be analyzed in terms of knowledge, indicating that he doesn’t take this approach to be an alternative to the Knowledge Analysis (which he explicitly endorses in DeRose 1991). As far as I’m aware, the only author who explicitly defends a Certainty Analysis of modals as an alternative to the Knowledge Analysis is Littlejohn (2011). On Littlejohn’s proposal, p is epistemically possible for S iff $\neg p$ is not obviously entailed by something S knows with certainty. While my proposal differs from Littlejohn’s in certain points of detail, in large part this section can be seen as making an extended case for Littlejohn’s thesis, and embedding this thesis within the broader framework of a Certainty First epistemology.

⁴⁷There are different ways of fleshing this out, depending on which standards of certainty we take to be in play. For those attracted to a contextualist account of modals, it’s natural to take the standards of certainty to be fixed by the context of utterance. More precisely:

$$\begin{aligned} \llbracket \Diamond \phi \rrbracket^{c,w} = 1 & \text{ iff } \cap \{p: \llbracket p \text{ is certain} \rrbracket^{c,w} = 1\} \cap \llbracket \phi \rrbracket^c \neq \emptyset \\ \llbracket \Box \phi \rrbracket^{c,w} = 1 & \text{ iff } \cap \{p: \llbracket p \text{ is certain} \rrbracket^{c,w} = 1\} \subseteq \llbracket \phi \rrbracket^c \end{aligned}$$

Relativists can take the standards of certainty to be fixed by the context of assessment.

(73) The butler must have done it.

We'd expect her to be also willing to assert:

(74) It's certain that the butler did it.

Indeed, it seems incoherent to assert (73) while denying (74):

(75) # The butler must have done it. But it's not certain that the butler did it.

The Certainty Analysis offers a straightforward explanation of this data. According to the Certainty Analysis, (73) says that it's epistemically certain that the butler did it (or that it's entailed by what's epistemically certain that the butler did it). And so (73) commits the detective to (74). For the same reason, (75) is predicted to be incoherent.

It's worth noting that we can reproduce the data using subjective certainty:

(76) # The butler must have done it. But I'm not certain that the butler did it.

One option for explaining the infelicity of (76) is to refine our analysis of epistemic modals, making them sensitive to both epistemic and subjective certainty (as in Littlejohn 2011). But another option is to retain our original analysis and explain the infelicity of (76) by appealing to the Matching Requirement. We expect that if p is epistemically certain for an agent, then that agent will be subjectively certain that p (as the Matching Requirement dictates). This expectation, together with our original Certainty Analysis, suffices to explain the bizarreness of (76).

Whereas the Certainty Analysis straightforwardly predicts the close connection between epistemic modals and certainty ascriptions, the Knowledge Analysis leaves this connection unexplained. According to the Knowledge Analysis, (73) says that the relevant folks' knowledge entails the butler did it. Since knowledge doesn't entail certainty, the fact that (73) is assertible is no guarantee that (74) is assertible. Similarly, we'd expect (75) to be coherent: it will be true whenever it is known but not certain that the butler did it.

Some might think this is too quick. One way proponents of the Knowledge Analysis might seek to explain these data is by appealing to the idea that we are reluctant to let the "contextually set standards for knowledge and certainty diverge" (Williamson 2000a: 204).

On this view, while knowledge doesn't entail either epistemic certainty or subjective certainty, a knowledge ascription will typically be true in a context c only if the corresponding epistemic and subjective certainty ascriptions are also true in c .

However, the same considerations that count against the idea that knowledge entails certainty cast doubt on the idea that we're reluctant to tolerate much separation between the standards for knowledge and the standards for certainty. Take, for example, the fact that *knows for certain* is not redundant: claiming that someone knows something with certainty is naturally interpreted as claiming that they know it with an unusually high degree of certainty (§3.2.1). This suggests that the standards for certainty ascriptions are generally higher than those for knowledge ascriptions. Moreover, there are cases where it seems natural to describe an agent as knowing a proposition without being certain of it. Consider the unconfident examinee, who reliably but hesitantly gives correct answers on an exam (Woozley 1952; Radford 1966). It seems that the examinee knows the answer to the examiner's question; however, it doesn't seem that the examinee knows the answer with either subjective or epistemic certainty (Armstrong 1969; Stanley 2008; McGlynn 2014: chp.5).

Alternatively, proponents of the Knowledge Analysis might try to explain the data by suggesting that while one can know p even though p isn't certain, knowing that p isn't certain precludes one from knowing that one knows p . If we assume that one can only felicitously assert what one is in a position to know, it follows that (75) is never assertable.

However, there are reasons for resisting an explanation along these lines. First, such an explanation would at most explain the oddity of (75). It wouldn't explain the intuition that whenever one is in a position to assert (73), one is also in a position to assert (74). Second, there remains a question as to *why* recognizing that p isn't certain would preclude one from knowing that one knows p . It seems that proponents of the Knowledge Analysis owe us a story here. Whatever story is offered, the resulting explanation of the infelicity of (75) will be less parsimonious than that offered by the Certainty Analysis, which treats (75) as a straightforward contradiction.

A further difficulty for pragmatic explanations of the data is that the connection between epistemic certainty and modals persists under embeddings. First, consider attitude verbs:

(77) ?? The detective thinks/knows the butler must have done it. But she doesn't think/know it's certain that he did it.

(78) ?? Suppose the butler must have done it but it's not certain whether he did it.

Second, consider conditionals:

(79) # If the butler must have done it and it's not certain whether he did it...

All of these sentences sound odd, if not downright incoherent. This provides further evidence that the connection between epistemic modals and epistemic certainty is not merely a pragmatic phenomenon, but calls for a semantic explanation. Again, a Certainty Analysis is in a better position than the Knowledge Analysis to explain the data.

For those who remain unconvinced, it is worth developing a final argument in favor of the Certainty Analysis—an argument from *modal concord*. Modal concord occurs when two modals occur near each other, but seem to only contribute the force of a single modal (Halliday 1970; Geurts and Huitink 2006; Huitink 2012). Consider:

(80) You may possibly have read my little monograph on the subject.⁴⁸

(81) You may have read my little monograph on the subject.

(80) is naturally interpreted as equivalent to (81) (or a slightly hedged version thereof). It is not naturally given a 'cumulative' reading, according to which it's possible that there is a possibility that the addressee has read the speaker's monograph.

It is a matter of debate how best to analyze modal concord. However, it's widely agreed that in order for modal concord to arise, the two modals must be equivalent. This explains why a concord reading is available for (80), but not for (82) or (83):

(82) ? Possibly you must have read my monograph.

(83) ? You may certainly have read my monograph.⁴⁹

⁴⁸Geurts and Huitink 2006 take (80) from Sir Arthur Conan Doyle's *The Hound of the Baskervilles*.

⁴⁹Arguably, we can access *some* coherent readings of these sentences. For instance, there seems to be a coherent reading of (83) on which it says that it is possible that it is certain that the addressee has read the speaker's monograph. However, this would be a cumulative reading rather than a concord reading.

This provides further reason to think that *must* and *certain* are equivalent. After all, *must* and *certainly* generate modal concord:

(84) You must certainly have read my monograph.

(85) You must have read my monograph.

(84) is naturally interpreted as equivalent to (85) (or a slightly strengthened version thereof). It is far less natural to give (84) a cumulative reading, according to which it is epistemically necessary that it is certain that the addressee read the monograph. Those who reject the Certainty Analysis will thus be forced to reject the well-supported generalization that modal concord is only possible when both modals are equivalent.⁵⁰

3.4.3 An Objection to the Certainty Analysis

I've argued that the Certainty Analysis accounts for a range of linguistic data that defenders of rival analyses—in particular, the Knowledge Analysis—will be hard-pressed to explain. However, some might object that the Certainty Analysis flounders when it comes to *epistemic contradictions* and *concessive knowledge attributions*:

(86) # It's raining but it might not be raining.

(87) # I know that it's raining. But it might not be.⁵¹

According to the Certainty Analysis, *It might not be raining* is true as long as it's not certain that it's raining. Since knowledge doesn't require certainty, it's not clear why a speaker couldn't both know that it's raining and also know that it's not certain (for her) that it's raining. And so it's unclear why (86) and (87) are infelicitous.

⁵⁰Geurts and Huitink 2006 suggest that modal concord may only require that the two modals are *nearly equivalent*. If this is right, then one could deny the Certainty Analysis while maintaining that the meaning of *must* is sufficiently close to the meaning of *certainly* to enable modal concord. Geurts and Huitink motivate their suggestion on the basis of sentences such as:

(iv) ? Pain in these diseases may probably influence the sleep process.

However, (iv) strikes me as odd. Insofar as I can interpret it, I find myself accessing a cumulative rather than a concord reading. (Informal polling suggests that a number of others share this judgment.) If this is right, then (iv) confirms rather than refutes the generalization that modal concord requires equivalence in meaning.

⁵¹Yalcin (2007) calls sentences such as (86), 'epistemic contradictions.' Rysiew (2001) calls discourses such as (87), 'concessive knowledge attributions.'

By contrast, the Knowledge Analysis seems to be in a comparatively good position to explain these data (Stanley 2005b). On the Knowledge Analysis, *It might not be raining* entails that the relevant agents don't know that it's raining. If we assume the 'Speaker Inclusion Constraint', according to which the speaker is always one of the relevant agents (Egan et al. 2005), (87) is contradictory.

While (86) is not contradictory on the Knowledge Analysis, proponents of the Knowledge Analysis arguably have a plausible pragmatic explanation of its infelicity. Recently, a number of authors have advocated the idea that knowledge is the norm of assertion:

Knowledge Norm of Assertion (KA): Assert p only if you know p .⁵²

By conjoining the Knowledge Analysis with KA, (87) is guaranteed to be unassertable. Since knowledge is factive, knowing the first conjunct (*It's raining*) precludes knowing the second conjunct (*It might not be*). On the face of it, this is an elegant result, since many have motivated KA on the grounds that it can explain the infelicity of Moorean assertions (Moore 1962), e.g.:

(88) # It's raining but I don't believe/know it's raining.

Hence by appealing to KA, proponents of the Knowledge Analysis can give a unified account of the infelicity of epistemic contradictions and Moorean assertions.

There are two ways that proponents of the Certainty Analysis could respond to this objection. The first is to replace KA with a certainty norm for assertion, e.g.:

Certainty Norm of Assertion (CA): Assert p only if p is epistemically certain for you.^{53,54}

⁵²See, in particular, Williamson 1996, 2000a. Unger (1975), Slote (1979), Reynolds (2002), DeRose (2002, 2009), Hawthorne (2004), Sutton (2005, 2007), Schaffer (2008), and Turri (2011a) all endorse versions of KA.

⁵³The idea for such a norm can be traced to Moore (1959), who claims that when I assert p , I imply that p is certain. However, Moore thought that knowledge entailed certainty. As far as I know, the only author who explicitly advocates a certainty norm as an alternative to KA is Stanley (2008). (I should note that since writing his 2008 paper, Stanley has expressed reservations about a certainty norm of assertion (p.c.), and it is doubtful whether he still endorses it.)

⁵⁴As Stanley (2008) observes, a norm along these lines seems most plausible if the standards of certainty are taken to be set by the asserter's context. On this construal, the norm says that you should only assert p in a context c if p counts as epistemically certain in c .

Armed with CA, advocates of the Certainty Analysis can give a pragmatic explanation of the oddity of (86). In order to felicitously utter the first conjunct (*It's raining*), it would need to be certain that it's raining. In order to felicitously utter the second conjunct (*It might not be raining*), it would need to be epistemically certain that it's not certain that it's raining. Since pos form epistemic certainty ascriptions are factive, a speaker could never satisfy both of these requirements. And since knowledge is factive, the same explanation generalizes to explain the infelicity of (87).

Some may worry that CA is ad hoc. But it can be motivated on independent grounds. As Unger (1975) and Stanley (2008) note, it sounds odd to say:

(89) ?? It's raining but it's not certain that it's raining.

What's more, such sentences seem to be defective in a way that closely resembles the original Moorean assertions (e.g., (88)). It would be nice to give a unified explanation of their oddity.

Given that knowledge doesn't entail epistemic certainty, KA doesn't explain the oddity of (89). CA does: it would be impossible to utter (89) while obeying CA. Hence (89) suffers from the same ailment as (86). Given the assumption that epistemic certainty entails knowledge (§3.2.5), CA also accounts for the original Moorean assertions that motivated KA.

It is worth mentioning a further advantage of CA. Given our analysis of *certain* as a max adjective (§3.2), CA also explains the oddity of:

(90) ? It's raining but it's not absolutely certain that it's raining.

According to CA, it's only felicitous to assert the first conjunct of (90) if the degree of epistemic certainty that it's raining is close enough to the max. If maximizing modifiers such as *absolutely* are also sensitive to standards of precision, then the second conjunct of (90) says that the degree of epistemic certainty that it's raining is less than whatever counts as close enough to the max. At the same time, we can explain why (90) sounds less bizarre than other Moorean assertions. After all, maximizing modifiers typically raise the standards of precision (§3.2.3). And this enables us to access a 'repair' interpretation of (90): perhaps the speaker is communicating that it is certain to rain, given low standards

of precision, but not given high standards.⁵⁵

An alternative strategy for explaining the infelicity of epistemic contradictions and concessive knowledge attributions is to modify the Certainty Analysis so as to predict that both (86) and (87) are semantically defective. For example, one could recast the Certainty Analysis as a version of *update semantics* (Veltman 1996). According to update semantics, the meaning of a sentence is a function from information states to information states, where an information state is a set of worlds compatible with the epistemic state of some relevant agent(s). An atomic sentence such as *It's raining* updates an information state s by removing any not-raining worlds from s . By contrast, sentences containing epistemic modals are tests on an information state. Asserting *It might not be raining* tests to see whether s contains at least one world where it isn't raining. If so, s passes the test, and is returned unscathed. If not, s crashes, returning the absurd information state (\emptyset). Similarly, asserting *It must be raining* tests whether s contains only worlds where it's raining.⁵⁶

Where does certainty come in? Formulations of update semantics rarely take a stand on what *sort* of epistemic state information states are supposed to model. We can recast the Certainty Analysis as an answer to this question: an information state just is the set of possibilities compatible with the propositions that are epistemically certain for the relevant agents. Note that this 'Updated Certainty Analysis' still accounts for the close connection between epistemic modals and certainty that motivated our original Certainty Analysis (§3.4.2).

⁵⁵This raises a subtle difference between the current proposal and Stanley's. While Stanley agrees that *certain* is an absolute gradable adjective, by this he seems to mean only that *certain* accepts the modifier *absolutely*; he denies that p is *certain* entails p is *absolutely certain* (Stanley 2008: 53-54). And so his account doesn't explain the oddity of (90).

⁵⁶To state this more precisely, let s be any information state, α be any atomic sentence, and ϕ and ψ be any sentences. Update semantics can be characterized by a function $[\cdot]$ from information states to information states, defined recursively as follows:

1. $s[\alpha] = s \cap \{w : w(\alpha) = 1\}$
2. $s[\phi \wedge \psi] = s[\phi][\psi]$
3. $s[\neg\phi] = s - s[\phi]$
4. $s[\Diamond\phi] = \{w \in s : s[\phi] \neq \emptyset\}$.
5. $s[\Box\phi] = \{w \in s : s \subseteq s[\phi]\}$.

For discussion and development of update semantics, see Gillies 2001; Yalcin 2012b, 2015; Willer 2013.

Recasting the Certainty Analysis in an update semantics framework enables a semantic diagnosis of the infelicity of epistemic contradictions and concessive knowledge attributions. Consider again (86) (*It's raining and it might not be raining*). Updating an information state with the first conjunct results in an information state that contains only worlds where it's raining. And so this information state is bound to fail the test imposed by the second conjunct. Therefore (86) comes out as semantically—not merely pragmatically—defective (Veltman 1996; Gillies 2001). Given the factivity of knowledge, this explanation generalizes to predict that (87) will also crash any information state.

Thus while the infelicity of epistemic contradictions and concessive knowledge attributions is a *prima facie* hurdle for the Certainty Analysis, there are two natural strategies for explaining the data: one pragmatic, one semantic. Which of these strategies is preferable? This is a difficult matter to adjudicate. On the one hand, we saw that a certainty norm for assertion can be motivated on independent grounds (specifically, its ability to explain the full range of Moorean assertions). This seems to be an important point in favor of the pragmatic strategy. On the other hand, Yalcin (2007) argues that epistemic contradictions are infelicitous in embedded contexts, unlike Moorean assertions. According to Yalcin, this creates a difficulty for purely pragmatic explanations of the infelicity of epistemic contradictions. If this is right, then this may count in favor of a semantic explanation of the data of the sort provided by the Updated Certainty Analysis.⁵⁷

For present purposes, we need not choose between the two strategies. (It may even turn out that both an update semantics for modals and a certainty norm for assertion are needed to account for the full range of data.) The important point is that these strategies offer ample resources for warding off the main objection to analyzing epistemic modals in terms of certainty.

⁵⁷More specifically, Yalcin argues that a pragmatic account of the infelicity of epistemic contradictions won't explain the oddity of:

- (v) ?? Suppose it's raining and it might not be raining.
- (vi) # If it's raining and might not be raining...

For a reply to Yalcin, see Dorr and Hawthorne 2013.

3.4.4 Taking Stock

Our ordinary uses of epistemic modals suggest that they're closely tied to certainty. This motivates a Certainty Analysis, according to which epistemic modals quantify over a set of possibilities whose domain is restricted by the propositions that are epistemically certain.

The Certainty Analysis also fits naturally with the treatment of evidential probability developed in §3.3. Evidential uses of probabilistic language seem to be closely related to epistemic modals. Earlier, we saw that *Probably* ϕ generates a scalar implicature that *It's not certain that* ϕ (§3.3.1). Similarly, probability claims implicate the falsity of the corresponding strong necessity claims:

- (91) a. It's likely the butler did it. \rightsquigarrow
 b. \neg (The butler must have done it).

And possibility claims implicate the falsity of the corresponding probability claims (Horn 1972):

- (92) a. The butler might have done it. \rightsquigarrow
 b. \neg (The butler probably did it).

On the picture that emerges, epistemic certainty ascriptions, epistemic modals, and evidential uses of probability operators all reside on the epistemic certainty scale. Pos form epistemic certainty ascriptions and strong necessity modals target the top of the scale: both are used to indicate that a proposition is maximally certain (or close enough thereto). Epistemic probability operators tend to live lower on the scale: they indicate that a proposition has a fairly high degree of epistemic certainty. Finally, epistemic possibility modals inhabit the bottom of scale: they indicate that a proposition isn't ruled out by what's epistemically certain.

3.5 Knowledge

Where does knowledge fit into this picture? Saying someone knows something with certainty seems to convey that she knows it on particularly strong grounds, or with a particularly

high degree of justification. (Recall again our minimal pairs from §3.2.1: (42a) vs. (42b); (43a) vs. (43b).) This suggests the following possibility: while knowledge requires a fairly high degree of epistemic certainty, the degree required isn't always as high as the degree required for the truth of a pos form epistemic certainty ascription. Perhaps knowledge is *near epistemic certainty*.⁵⁸

According to the account developed in §3.2, ' p is certain for S ' is true as long as p 's degree of epistemic certainty is close enough to the maximum degree of epistemic certainty. Thus saying that knowledge is near certainty amounts to the following:

Close to Close: ' S knows p ' is true in c iff p 's degree of epistemic certainty (for S) is close enough to being close enough (for the purposes of c) to the maximum.

Here's another way of putting it. According to the treatment in §3.2, there will be an interval of degrees of epistemic certainty that are indistinguishable from the maximum for the purposes of the context. Any difference between them and the maximum is ignorable. On the present proposal, there will often be another interval adjacent to the first. This is the range of *tolerable* differences: they're not indistinguishable from the max, but they're nearly indistinguishable. ' S knows p ' is true iff p 's degree of epistemic certainty (for S) is at least tolerable.

How does this analysis handle false beliefs and Gettier cases? In the same way that our treatment of pos form epistemic certainty ascriptions did. Recall that I proposed explaining the puzzling entailments of pos form epistemic certainty ascriptions by imposing substantive constraints on the 'close enough' relation (§3.2.5). For example, to explain why pos form epistemic certainty ascriptions are factive, we imposed a Truth Constraint: if p is false, then no degree of epistemic certainty less than 1 is ignorable. We can likewise impose a Truth Constraint on the 'close to close' relation: if p is false, then no degree of epistemic certainty less than 1 is tolerable.

Similarly, for any anti-Gettier condition C , we can regard C as a constraint on both the 'close enough' relation and the 'close to close' relation. For example, those attracted to

⁵⁸Obviously, this slogan needs to be understood in such a way that epistemic certainty trivially counts as near epistemic certainty.

the idea that a safety condition on knowledge explains why Gettier subjects lack knowledge could impose a Safety Constraint on the ‘close to close’ relation: if S’s belief that p is unsafe, then no degree of epistemic certainty less than 1 is tolerable. (See Figure 3.2.)

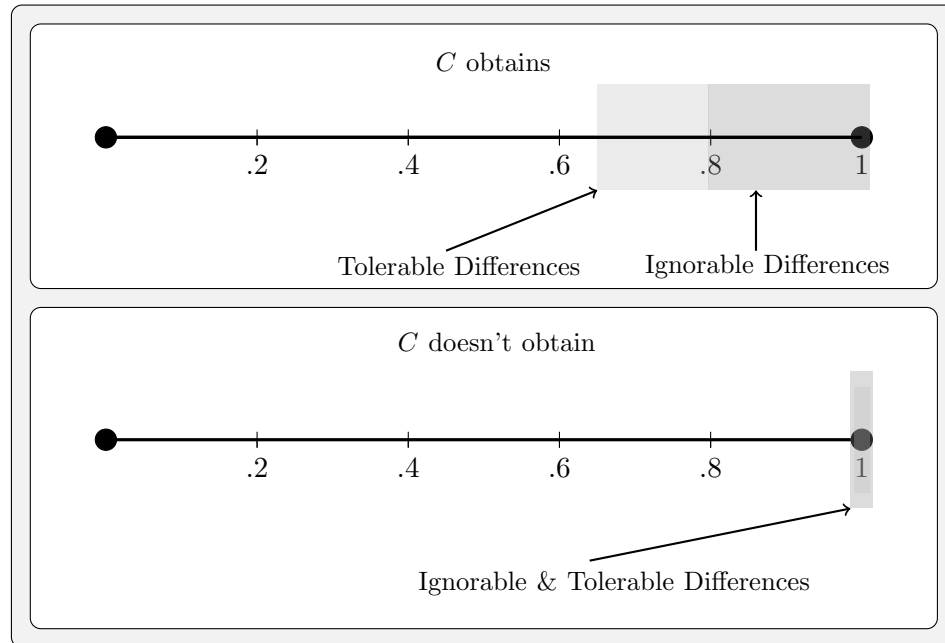


Figure 3.2: Certainty & Near Certainty.

What if it turns out that there’s no anti-Gettier condition? This would mean that there’s no non-circular specification of the conditions under which a degree of epistemic certainty counts as close to close enough to the max. But even if this turns out to be the case, the present account is still illuminating. While we won’t have given a reductive analysis of knowledge, we’ll still have shed light on the structural relations between knowledge and certainty: to know p is for p ’s degree of epistemic certainty to come close to that which suffices for certainty (full stop). This is an interesting, substantive connection—one that does not emerge from rival accounts.

3.6 Certainty and Action

Thus far, I’ve argued that certainty can be used to shed light on evidential probability, epistemic modals, and knowledge. I’ve also suggested that some important normative constraints are best articulated in terms of certainty. In particular, I’ve argued that our degrees

of subjective certainty ought to align with the degrees of epistemic certainty (the Matching Requirement); I've also highlighted the attractions of a certainty norm of assertion (§3.4.3). In this section, I explore whether we can use certainty to provide normative constraints on action.

One attractive picture of the normative constraints on action is provided by Bayesian decision theory: agents ought to maximize expected utility. What probability function should be used when calculating expected utilities? A natural option is to employ an agent's evidential probability function (Hawthorne 2005; Greco 2013). Given our earlier arguments for interpreting evidential probabilities as degrees of epistemic certainty (§3.3), this suggests one role for certainty in constraining action: degrees of certainty feature in the interpretation of expected utilities.⁵⁹

Recently, a number of authors have suggested that while Bayesian decision theory may be correct, it does not provide the complete story about the normative constraints on action. In the real world, we often take some premises for granted. In deciding when to leave my apartment, I do not typically calculate the probability that my car will start; instead, I simply rely on the premise that my car will start. This raises the question: what sort of epistemic relation does one need to stand in towards a premise in order to appropriately rely on it in practical reasoning?

According to many authors in the recent literature, the answer is: 'Knowledge.'⁶⁰ Here's a simple way of formulating this idea:

Knowledge Norm of Action (KN): It is appropriate for S to rely on p in practical reasoning iff S knows p .

KN holds considerable *prima facie* appeal. As Hawthorne and Stanley note, we often criticize others for acting on the basis of propositions that they don't know: "If a parent

⁵⁹Greco (2013) argues that if we (i) calculate expected utility using the evidential probability function, (ii) adopt the Knowledge Account of Evidential Probability, we'll be led to recommend accepting irrational bets—e.g., bets where I gain a penny if *Snow* (the proposition that it will snow in NY sometime this year) is true, and where the world is destroyed if *Snow* is false. This reveals a further advantage of taking evidential probabilities to be degrees of certainty: no such bets are recommended, because even though I know *Snow*, its degree of epistemic certainty—and hence its evidential probability for me—is less than 1.

⁶⁰While they differ on points of detail, Hawthorne (2004); Williamson (2005); Fantl and McGrath (2002, 2009); Hawthorne and Stanley (2008); Weatherston (2012); and Ross and Schroeder (2014) all defend the idea that knowledge places an important constraint on action.

allows a child to play near a dog and does not know whether the dog would bite the child, and if a doctor uses a needle that he did not know to be safe, then they are *prima facie* negligent” (2008: 572).

But while KN explains such criticisms, it doesn’t explain the fact that we’re often willing to criticize agents for acting on the basis of uncertain premises (Brown 2008a,b, 2012b; Gerken 2011):

(93) You shouldn’t let your child play there, since we don’t know for sure whether that dog is friendly.

(94) You shouldn’t use that needle, since it’s not certain whether it’s safe.

Confronted with the naturalness of criticizing people for their lack of certainty, one might conclude that we shouldn’t lean too heavily on our ordinary patterns of criticism (Gerken 2011), or that the epistemic state required for appropriate action varies with circumstances (Brown 2008a,b, 2012b). But another option is to explain these criticisms via a certainty norm:

Certainty Norm of Action (CN): It is appropriate for S to rely on p in practical reasoning iff p is certain for S.⁶¹

CN straightforwardly explains the naturalness of criticizing actions on the basis of uncertain premises. Assuming that epistemic certainty entails knowledge, CN also explains the naturalness of criticizing people for acting on the basis of unknown propositions. Thus CN holds out the promise of a general explanation of our ordinary patterns of criticism.⁶²

A separate argument for CN starts from the question: ‘What is to rely on a proposition in practical reasoning?’ It seems plausible that when I rely on p in practical reasoning, I treat any $\neg p$ possibilities as ignorable, at least for the purposes of deliberation. (If I didn’t treat such possibilities as ignorable, then it seems I wouldn’t be relying on p ; instead, I

⁶¹As in the case of a Certainty Norm of Assertion, a Certainty Norm of Action is most plausible if the standards for certainty are provided by S’s context.

⁶²CN is thus a factive (and, indeed, knowledge-entailing) norm of action. In this regard, it is importantly different from non-factive norms of action, for instance, the sort of variable justification/warrant norm defended by Gerken (2011, forthcoming) and Locke (2015), which holds (roughly) that it is appropriate to act on the basis of p iff S is justified/warranted in believing p to the degree that is required by the circumstances.

would be relying on a proposition of the form: $\langle \text{It's very likely that } p \rangle$.) This conception of what it is to rely on a proposition in practical reasoning in turn motivates the following idea: it's appropriate for S to rely on p in a context c iff it's appropriate for S to ignore any $\neg p$ possibilities in c .⁶³ And it seems very natural to hold that it's appropriate for S to ignore any $\neg p$ possibilities in c iff any difference between p 's degree of certainty (for S) and the maximum can be appropriately ignored in c . Given Close Enough (§3.2.4), this obtains iff 'It is certain that p ' is true in c .

CN also offers to handle some of the counterexamples to KN in the literature. Consider, for instance, Brown (2008a)'s surgeon case:

A student observes a surgeon examining patient A who has a diseased left kidney. The decision is taken to remove it that afternoon. Later, the student observes the surgeon preparing to operate. As patient A lies on the table, the surgeon consults the patient's notes. The student turns to one of the nurses:

Student: I don't understand. Why is she looking at the patient's records? She was in clinic with the patient this morning. Doesn't she even know which kidney it is?

Nurse: Of course she knows which kidney it is. But, imagine what it would be like if she removed the wrong kidney. She shouldn't operate before checking the patient's records.⁶⁴

CN explains the nurse's appraisal. Even though the surgeon knows that the left kidney is diseased, she shouldn't rely on this proposition, since it isn't epistemically certain for her: its degree of certainty isn't close enough to the maximum.

According to this diagnosis of Brown's surgeon case, practical factors can affect whether p 's degree of epistemic certainty is close enough to the maximum in a given context. (It's the disastrous consequences of removing the wrong kidney which ratchet up the standards for certainty in the surgeon's context of deliberation.) Thus we are led to a sort of 'impurism' about certainty ascriptions, according to which the truth-value of 'It's certain that p ' in a

⁶³Or to frame this in the terms suggested by Weatherson (2012), one can appropriately rely on p iff one can appropriately leave any $\neg p$ possibilities off of a decision table.

⁶⁴See Lackey 2010 and Reed 2010 for other counterexamples to the sufficiency direction of KN.

context partially depends on the practical consequences of being wrong about p . (Similar remarks seem to apply to pos form subjective certainty ascriptions.)⁶⁵ However, this does not seem to be a special feature about certainty *per se*: in general, practical factors influence the resolution of gradable adjectives. Consider, for example, *clean*. Whether a knife's degree of cleanness counts as close enough to the max depends on how important it is that the knife is completely free of dirt: surgery calls for higher standards of precision than carrot-chopping. Arguably, then, impurism about certainty ascriptions is a less surprising doctrine than impurism about knowledge ascriptions. The general observation that practical features affect the standards of precision for gradable expressions leads us to expect impurism about the former, but not the latter (since, after all, *knows* isn't gradable).⁶⁶

In summary, the hypothesis that certainty has a close connection with action holds considerable appeal. This connection emerges at two levels. It first emerges in our interpretation of Bayesian decision theory: rational agents maximize expected utility, and the probability function used to determine expected utilities is plausibly identified with the epistemic certainty function. The connection emerges a second time when it comes to articulating the conditions under which it's appropriate to rely on a proposition in practical reasoning. A certainty norm (CN) provides a unified account of our patterns of criticism; it meshes naturally with the view that relying on p involves ignoring any $\neg p$ possibilities; it also accommodates some of the counterexamples to knowledge norms of action.⁶⁷

Is there any reason to resist CN? Perhaps the most natural concern is that CN imposes too demanding a standard on practical reasoning (Weatherson 2012: 80-81). However, given

⁶⁵For defenses of impurism about knowledge ascriptions, see Fantl and McGrath 2002, 2009; Hawthorne 2004; Stanley 2005a; Weatherson 2011, 2012.

⁶⁶It's worth noting that if we replace KN with CN, we undercut one popular argument for impurism about knowledge. The argument in question starts with the observation that whether you can appropriately act on the basis of p varies with practical factors. It then appeals to KN in order to derive the conclusion that knowledge varies with practical factors. (For versions of this argument, see Stanley 2005a; Fantl and McGrath 2009; Weatherson 2012.) That said, CN is certainly compatible with impurism about knowledge. And impurism about knowledge fits fairly naturally with the picture of knowledge developed in §3.5: if the 'close enough' relation is affected by practical factors, it seems natural to think that the 'close to close enough' relation is affected by practical factors as well.

⁶⁷Arguably, CN gains further appeal in light of our earlier arguments for a certainty norm of assertion (§3.4.3). Of course, it is controversial whether the norm of assertion is derivable from the norm of action. (See Brown 2012a; Montminy 2013; McGlynn 2014: chp.6.) But even if we can't derive the former from the latter, the view that a single epistemic state imposes normative constraints on both assertion and action is arguably more elegant than a disjunctive picture, according to which different epistemic states govern the two.

the context-sensitive account of certainty advanced here, this worry loses much of its bite. Since much of our everyday knowledge counts—in ordinary contexts—as certain, much of our everyday knowledge will serve—again, in ordinary contexts—as an appropriate basis for action. And even when some ordinary item of knowledge p does not meet the contextual standards for certainty, this need not render us paralyzed: we can still use p 's high degree of certainty to calculate the expected utilities of various actions, as recommended by decision theory.⁶⁸

3.7 Conclusion

In this chapter, I've argued that certainty should occupy a central role in epistemology. I began by disputing the common view that certainty is scarce. According to the context-sensitive semantics for certainty ascriptions developed here, many of our beliefs qualify—in appropriate contexts—as both subjectively and epistemically certain. Having laid the foundation, I went on to put certainty to work in epistemology. I suggested that a variety of topics—evidential probability, epistemic modals, knowledge, and the normative constraints on credence, assertion, and action—can be fruitfully understood in terms of certainty.

Of course, there are a number of epistemological topics I have not broached in this chapter. It is natural to wonder which of these can also be understood in terms of certainty. Take, for example, the nature of belief. Knowledge Firsters often try to explain belief in terms of knowledge. According to a simple version of a Knowledge First treatment of belief, to believe p is to treat p as though you know it: merely believing is a sort of “botched knowing” (Williamson 2000a: 47).⁶⁹ From the present perspective, it's natural to trade this knowledge-centered analysis for an analysis in terms of degrees of subjective certainty. Just as knowledge is near epistemic certainty, belief is near subjective certainty: one believes p iff one's degree of subjective certainty in p is close enough to being close enough to the maximum. In slogan form: epistemic certainty is to knowledge as subjective certainty is to belief. I hope to explore the viability of this approach in future work.⁷⁰

⁶⁸Compare with Hawthorne and Stanley 2008, who argue that we often act on our knowledge of the evidential probabilities of various propositions.

⁶⁹For a similar account, see Adler 2002: 36, 275. For objections, see McGlynn 2013.

⁷⁰One interesting consequence of this proposal, when combined with CN, is that appropriately believing p

It is too early to tell whether a Certainty First program—or indeed any ‘X First’ program—should be embraced in full generality. However, I hope that this chapter has given grounds for cautious optimism. At the very least, I hope to have shown that certainty enjoys a relatively high degree of explanatory priority: we can make considerable progress by taking certainty as our primitive and using it explain other epistemological phenomena.

does not put one in a position to appropriately act on the basis of p . After all, if believing p only requires *near* subjective certainty, then one can rationally believe p even though p isn’t quite certain. This consequence seems plausible in light of discourses such as the following: ‘I believe our plane leaves at seven, but we had better check.’ In making this utterance, the speaker avows a belief that the plane leaves at seven, while at the same time denying that it would be appropriate to act on the basis of this proposition. (For an independent and complementary argument that appropriately believing p does not put one in a position to appropriately assert p , see Hawthorne et al. 2016.)

Bibliography

- Jonathan Adler. *Belief's Own Ethics*. MIT Press, Cambridge, MA, 2002.
- William Alston. The deontological conception of epistemic justification. *Philosophical Perspectives*, 2:115–152, 1988.
- Thomas Aquinas. *Commentary on the Posterior Analytics of Aristotle*. Magi Books Inc., New York, 1970.
- D.M. Armstrong. Does knowledge entail belief? *Proceedings of the Aristotelian Society*, 70: 21–36, 1969.
- A.J. Ayer. *Language, Truth, and Logic*. Dover, New York, 1936.
- Kelly Becker. Epistemic luck and the generality problem. *Philosophical Studies*, 139:353–366, 2008.
- Bob Beddor. Evidentialism, circularity, and grounding. *Philosophical Studies*, 172:1847–1868, 2015a.
- Bob Beddor. Process reliabilism's troubles with defeat. *The Philosophical Quarterly*, 65 (259), 2015b.
- Bob Beddor. Justification as faultlessness. *Philosophical Studies*, forthcoming.
- Matthew Bedke. Developmental process reliabilism: on justification, defeat, and evidence. *Erkenntnis*, 73(1):1–17, 2010.
- Gunnar Björnsson and Tristram McPherson. Moral attitudes for non-cognitivists: Solving the specification problem. *Mind*, 123(408):1–38, 2014.
- Deborah Black. Knowledge (‘ilm) and certitude (yaqin). *Arabic Sciences and Philosophy*, 16(1):11–45, 2006.
- Simon Blackburn. *Spreading the Word*. Clarendon Press, Oxford, 1984.

- Simon Blackburn. *Ruling Passions*. Clarendon Press, Oxford, 1998.
- Tomas Bogardus. Knowledge under threat. *Philosophy and Phenomenological Research*, 86(1):289–313, 2014.
- Laurence Bonjour. *The Structure of Empirical Knowledge*. Harvard University Press, Cambridge, MA, 1985.
- Berit Brogaard. Moral contextualism and moral relativism. *Philosophical Quarterly*, 58:385–409, 2008.
- Jessica Brown. Subject-sensitive invariantism and the knowledge norm for practical reasoning. *Noûs*, 42(2):167–189, 2008a.
- Jessica Brown. Knowledge and practical reason. *Philosophy Compass*, 3(6):1135–1152, 2008b.
- Jessica Brown. Assertion and practical reasoning: Common or divergent standards? *Philosophy and Phenomenological Research*, 84(1):123–157, 2012a.
- Jessica Brown. Practical reasoning, decision theory and anti-intellectualism. *Episteme*, 9(1):43–62, 2012b.
- Herman Cappelen and John Hawthorne. *Relativism and Monadic Truth*. Oxford University Press, Oxford, 2009.
- Nate Charlow. ‘prospects for an expressivist theory of meaning’. *Philosophers’ Imprint*, 15(23), 2015.
- Matthew Chrisman. From epistemic contextualism to epistemic expressivism. *Philosophical Studies*, 135(2):225–254, 2007.
- Matthew Chrisman. Epistemic expressivism. *Philosophy Compass*, 7(2):118–126, 2012a.
- Matthew Chrisman. ‘ought’ and control. *Australasian Journal of Philosophy*, 90(3):433–451, 2012b.
- Matthew Chrisman. On the meaning of ‘ought’. In Russ Shafer-Landau, editor, *Oxford Studies in Metaethics*, volume 7. Oxford University Press, New York, 2012c.
- Philippe Chuard and Nicholas Southwood. Epistemic norms without voluntary control. *Noûs*, 43(4):599–632, 2009.
- Michael J. Clark and David Liggins. Recent work on grounding. *Analysis*, 72(4):812–823, 2012.

- Juan Comesaña. Unsafe knowledge. *Synthese*, 146:395–404, 2005.
- Juan Comesaña. Evidentialist reliabilism. *Noûs*, 44(4):571–600, 2010.
- Earl Conee and Richard Feldman. The generality problem for reliabilism. *Philosophical Studies*, 89(1):1–29, 1998.
- Earl Conee and Richard Feldman. *Evidentialism: Essays in Epistemology*. Oxford University Press, New York, 2004.
- Earl Conee and Richard Feldman. Some virtues of evidentialism. *Veritas*, 50:95–108, 2005.
- Earl Conee and Richard Feldman. Evidence. In Smith, editor, *Epistemology: New Essays*. Oxford University Press, New York, 2008.
- Bridget Copley. What should *should* mean?, 2006. URL <http://copley.free.fr/copley.should.pdf>.
- M.J. Cresswell. The semantics of degree. In Barbara Partee, editor, *Montague Grammar*. Academic Press, New York, 1977.
- López de Sa. Presuppositions of commonality: An indexical relativist account of disagreement. In García Carpintero and Kölbel, editors, *Relative Truth*. Oxford University Press, Oxford, 2008.
- Keith DeRose. Epistemic possibility. *Philosophical Review*, 100:581–605, 1991.
- Keith DeRose. Simple ‘might’s, indicative possibilities, and the open future. *Philosophical Quarterly*, 48:67–82, 1998.
- Keith DeRose. Assertion, knowledge, and context. *Philosophical Review*, 111(2):167–203, 2002.
- Keith DeRose. *The Case for Contextualism: Knowledge, Skepticism, and Context*. Clarendon Press, Oxford, 2009.
- Rene Descartes. *Rules for the Direction of the Mind*. In Adam and Tannery, editors, *Oeuvres De Descartes*. Librairie Philosophique, Paris, 1983.
- Sinan Dogramaci. Reverse engineering epistemic evaluations. *Philosophy and Phenomenological Research*, 84(3):513–530, 2012.
- Zoltan Domotor. Probability kinematics and representation of belief change. *Philosophy of Science*, 47, 1980.
- Cian Dorr and John Hawthorne. Embedding epistemic modals. *Mind*, 2013.

- Jamie Dreier. Internalism and speaker relativism. *Ethics*, 101(1):6–26, 1990.
- Andy Egan. Epistemic modals, relativism, and assertion. *Philosophical Studies*, 133(1): 1–22, 2007.
- Andy Egan. Disputing about taste. In Feldman and Warfield, editors, *Disagreement*. Oxford University Press, Oxford, 2010.
- Andy Egan. Relativist dispositional theories of value. *Southern Journal of Philosophy*, 50 (4):557–582, 2012.
- Andy Egan. There’s something funny about comedy: A case study in faultless disagreement. *Erkenntnis*, 79(1):73–100, 2014.
- Andy Egan and Brian Weatherson. Epistemic modals and epistemic modality. In Egan and Weatherson, editors, *Epistemic Modality*. Oxford University Press, Oxford, 2011.
- Andy Egan, John Hawthorne, and Brian Weatherson. Epistemic modals in context. In Preyer and Peter, editors, *Contextualism in Philosophy*. Oxford University Press, Oxford, 2005.
- Jeremy Fantl and Matthew McGrath. Evidence, pragmatics, and justification. *Philosophical Review*, 111(1):67–94, 2002.
- Jeremy Fantl and Matthew McGrath. *Knowledge in an Uncertain World*. Oxford University Press, Oxford, 2009.
- Richard Feldman. Having evidence. In Austin, editor, *Essays Presented to Edmund Gettier*. Kluwer Academic Publishers, Dordrecht, 1988.
- Hartry Field. Epistemology without metaphysics. *Philosophical Studies*, 143(2):249–290, 2009.
- Kit Fine. The question of realism. *Philosophers’ Imprint*, 1(1):1–30, 2001.
- Kit Fine. Guide to ground. In Correia and Schneider, editors, *Metaphysical Grounding*. Cambridge University press, Cambridge, 2012.
- Jerry Fodor. *The Modularity of Mind*. MIT Press, 1983.
- Francesca Foppolo and Francesca Panzeri. Do children know when their room counts as clean? MS, University of Milano-Bicocca, 2010.
- Daniel Garber. Old evidence and logical omniscience in bayesian confirmation theory. In

- Earman, editor, *Testing Scientific Theories*. University of Minnesota Press, Minneapolis, 1983.
- Mikkel Gerken. Warrant and action. *Synthese*, 178(3):529–547, 2011.
- Mikkel Gerken. Against knowledge first epistemology. In Carter and Gordon, editors, *Knowledge First Epistemology*. Oxford University Press, Oxford, forthcoming.
- Bart Geurts and Janneke Huitink. Modal concord. In *Concord Phenomena and the Syntax-Semantics Interface*. ESSLLI, 2006.
- Allan Gibbard. *Wise Choices, Apt Feelings*. Harvard University Press, Cambridge, MA, 1990.
- Allan Gibbard. *Thinking How to Live*. Harvard University Press, Cambridge, MA, 2003.
- Allan Gibbard. Rational credence and the value of truth. In Gendler and Hawthorne, editors, *Oxford Studies in Epistemology*, volume 2. Oxford University Press, New York, 2007.
- Allan Gibbard. *Reconciling Our Aims*. Oxford University Press, Oxford, 2008.
- Thony Gillies. A new solution to moore’s paradox. *Philosophical Studies*, 105:237–250, 2001.
- Alvin Goldman. What is justified belief? In Pappas, editor, *Justification and Knowledge*. Reidel, Dordrecht, 1979.
- Alvin Goldman. *Epistemology and Cognition*. Harvard University Press, Cambridge, MA, 1986.
- Alvin Goldman. Epistemic folkways and scientific epistemology. In *Liasons: Philosophy Meets the Cognitive and Social Sciences*. MIT Press, Cambridge, MA, 1992.
- Alvin Goldman. *Knowledge in a Social World*. Oxford University Press, New York, 1999.
- Alvin Goldman. Epistemic relativism and reasonable disagreement. In Feldman and Warfield, editors, *Disagreement*, pages 187–215. Oxford University Press, New York, 2010.
- Alvin Goldman. Towards a synthesis of reliabilism and evidentialism. In Dougherty, editor, *Evidentialism and its Discontents*, pages 254–280. Oxford University Press, New York, 2011.
- Alvin Goldman and Moshe Shaked. An economic model of scientific activity and truth acquisition. *Philosophical Studies*, 63(1):31–55, 1991.

- Martin Grajner. Hybrid expressivism and epistemic justification. *Philosophical Studies*, 172(9):2349–2369, 2015.
- Daniel Greco. Prodigality and probability. In Gendler and Hawthorne, editors, *Oxford Studies in Epistemology*, volume 4, pages 82–107. Oxford University Press, Oxford, 2013.
- Daniel Greco. A puzzle about epistemic akrasia. *Philosophical Studies*, 167(2):201–219, 2014.
- Thomas Grundmann. Reliabilism and the problem of defeaters. *Grazer Philosophische Studien*, 79(1):65–76, 2009.
- Ian Hacking. Possibility. *Philosophical Review*, 76:143–168, 1967.
- Michael Alexander Kirkwood Halliday. Functional diversity in language as seen from a consideration of modality and mood in english. *Foundations of Language*, 6:322–261, 1970.
- Elizabeth Harman. Morally permissible moral mistakes. *Ethics*, 126(2):366–393, 2016.
- Gilbert Harman. Moral relativism defended. *Philosophical Review*, 84:3–22, 1975.
- Gilbert Harman. Stringency of rights and 'ought'. *Philosophy and Phenomenological Research*, 53:181–185, 1993.
- John Hawthorne. *Knowledge and Lotteries*. Oxford University Press, New York, 2004.
- John Hawthorne. Knowledge and evidence. *Philosophy and Phenomenological Research*, 70(2):452–458, 2005.
- John Hawthorne. Eavesdroppers and epistemic modals. *Philosophical Issues*, 17:92–101, 2007.
- John Hawthorne. Knowledge and epistemic necessity. *Philosophical Studies*, 158(3):493–501, 2012.
- John Hawthorne and Jason Stanley. Knowledge and action. *Journal of Philosophy*, 105(10):571–590, 2008.
- John Hawthorne, Daniel Rothschild, and Levi Spectre. Belief is weak. *Philosophical Studies*, 173:1393–1404, 2016.
- Trevor Hedberg. Epistemic supererogation and its implications. *Synthese*, 2014.
- Brian Hedden. Time-slice rationality. *Mind*, 124(494):449–491, 2015.
- Jaakko Hintikka. *Knowledge and Belief*. Cornell University Press, 1962.

- Terry Horgan and Mark Timmons. Cognitivist expressivism. In Horgan and Timmons, editors, *Metaethics After Moore*. Oxford University Press, Oxford, 2006.
- Laurence Horn. *On the Semantic Properties of the Logical Operators in English*. UCLA PhD dissertation, 1972.
- Laurence Horn. *A Natural History of Negation*. University of Chicago Press, Chicago, 1989.
- Michael Huemer. Compassionate phenomenal conservatism. *Philosophy and Phenomenological Research*, 74:30–55, 2007.
- Janneke Huitink. Modal concord: A case study in dutch. *Journal of Semantics*, 29(3): 403–437, 2012.
- Susan Hurley. *Natural Reasons*. Oxford University Press, Oxford, 1989.
- Richard Jeffrey. *The Logic of Decision*. University of Chicago Press, Chicago, 1965.
- Martin Jönsson. A reliabilism built on cognitive convergence: An empirically grounded solution to the generality problem. *Episteme*, 10(3):241–268, 2013a.
- Martin Jönsson. Linguistic convergence in verbs for belief-forming processes. *Philosophical Psychology*, 28(1):114–138, 2013b.
- James Joyce. A nonpragmatic vindication of probabilism. *Philosophy of Science*, 65(4): 575–603, 1998.
- James Joyce. Williamson on evidence and knowledge. *Philosophical Books*, 45(4):296–305, 2004.
- Mark Kaplan. Who cares what you know? *Philosophical Quarterly*, 53:105–116, 2003.
- Mark Kaplan. Williamson’s casual approach to probabilism. In Greenough and Pritchard, editors, *Williamson on Knowledge*. Oxford University Press, New York, 2009.
- Klemens Kappel. Expressivism about knowledge and the value of knowledge. *Acta Analytica*, 25(2):175–194, 2010.
- Antti Kauppinen. What makes a sentiment moral? In Shafer-Landau, editor, *Oxford Studies in Metaethics*, volume 1. Oxford University Press, New York, 2010.
- Christoph Kelp. Knowledge and safety. *Journal of Philosophical Research*, 34:21–31, 2009.
- Chris Kennedy. The semantics of relative and absolute gradable adjectives. *Linguistics and Philosophy*, 30:1–45, 2007.

- Chris Kennedy and Louise McNally. Scale structure, degree modification, and the semantics of gradable predicates. *Language*, 81:345–381, 2005.
- Jaegwon Kim. What is ‘naturalized epistemology’? *Philosophical Perspectives*, 2:381–405, 1988.
- Kihyeon Kim. The deontological conception of epistemic justification and doxastic voluntarism. *Analysis*, 54(4):282–284, 1994.
- Peter Klecha. Positive and conditional semantics for gradable modals. In Guevara et al, editor, *Proceedings of Sinn und Bedeutung 16*, volume 2, pages 363–376. MIT Working Papers in Linguistics, 2012.
- Peter Klein. *Certainty: A Refutation of Scepticism*. University of Minnesota Press, Minneapolis, 1981.
- Sebastian Köhler. Do expressivists have an attitude problem? *Ethics*, 123(3):479–507, 2013.
- Niko Kolodny and John MacFarlane. Ifs and oughts. *Journal of Philosophy*, 107(3):115–143, 2010.
- Hilary Kornblith. Justified belief and epistemically responsible action. *The Philosophical Review*, 92(1):33–48, 1983.
- Christine Korsgaard. *Creating the Kingdom of Ends*. Cambridge University Press, Cambridge, 1996.
- Angelika Kratzer. The notional category of modality. In Eikmeyer and Rieser, editors, *Words, Worlds, and Contexts: New Approaches in Word Semantics*. W. de Gruyter, Berlin, 1981.
- Angelika Kratzer. Modality. In von Stechow and Wunderlich, editors, *Semantics: An International Handbook of Contemporary Research*. W. de Gruyter, Berlin, 1991.
- Angelika Kratzer. *Modals and Conditionals*. Oxford University Press, Oxford, 2012.
- Manfred Krifka. Approximate interpretations of number words: A case for strategic communication. In Bouma, Kramer, and Zwarts, editors, *Cognitive Foundations of Communication*. Koninklijke Nederlandse Akademie van Wetenschappen, Berlin, 2007.
- Jonathan Kvanvig. Two approaches to epistemic defeat. In Baker, editor, *Alvin Plantinga*. Cambridge University Press, Cambridge, 2007.
- Jennifer Lackey. Acting on knowledge. *Philosophical Perspectives*, 24:361–382, 2010.

- Marc Lange. Is jeffrey conditionalization defective by virute of being non-commutative? remarks on the sameness of sensory experience. *Synthese*, 123, 2000.
- Peter Lasersohn. Pragmatic halos. *Language*, 75:522–551, 1999.
- Peter Lasersohn. Context dependence, disagreement, and predicates of personal taste. *Linguistics and Philosophy*, 28:643–686, 2005.
- Maria Lasonen-Aarnio. Unreasonable knowledge. *Philosophical Perspectives*, 24:1–21, 2010.
- Daniel Lassiter. Gradable epistemic modals, probability, and scale structure. In Li and Lutz, editors, *Semantics and Linguistic Theory (SALT)*, volume 20, pages 197–215. CLC Publications, Ithaca, New York, 2010.
- Daniel Lassiter. Measurement and modality: The scalar basis of modal semantics, 2011.
- Daniel Lassiter and Noah Goodman. How many kinds of reasoning? inference, probability, and natural language semantics. *Cognition*, 136:123–134, 2015.
- Keith Lehrer. *Theory of Knowledge*. Routledge, New York, 1990.
- Hannes Leitgeb. The stability theory of belief. *Philosophical Review*, 123(3):131–171, 2014.
- Isaac Levi. Probability kinematics. *British Journal for the Philosophy of Science*, 18: 197–209, 1967.
- David Lewis. Psychophysical and theoretical identifications. *Australasian Journal of Philosophy*, 50:249–258, 1972.
- David Lewis. Attitudes *De Dicto* and *De Se*. *Philosophical Review*, 88:513–543, 1979a.
- David Lewis. Scorekeeping in a language game. *Journal of Philosophical Logic*, 8(1):339–359, 1979b.
- David Lewis. Dispositional theories of value. *Proceedings of the Aristotelian Society*, 63, 1989.
- David Lewis. Reduction of mind. In Guttenplan, editor, *A Companion to Philosophy of Mind*. Blackwell, Oxford, 1994.
- David Lewis. Elusive knowledge. *Australasian Journal of Philosophy*, 74(4):549–567, 1996.
- Clayton Littlejohn. Concessive knowledge attributions and fallibilism. *Philosophy and Phenomenological Research*, 83(3):603–619, 2011.
- Clayton Littlejohn. *Justification and the Truth Connection*. Cambridge University Press, Cambridge, 2012.

- Clayton Littlejohn. How and why knowledge is first. In Carter and Gordon, editors, *Knowledge First Epistemology*. Oxford University Press, Oxford, forthcoming.
- Dustin Locke. Practical certainty. *Philosophy and Phenomenological Research*, 90(1):72–95, 2015.
- Jack Lyons. *Perception and Basic Beliefs*. Oxford University Press, New York, 2009.
- Jack Lyons. Goldman on evidence. In Kornblith and McLaughlin, editors, *Goldman and his Critics*. Blackwell, Oxford, forthcoming.
- John MacFarlane. Epistemic modals are assessment sensitive. In Egan and Weatherson, editors, *Epistemic Modality*. Oxford University Press, Oxford, 2011.
- John MacFarlane. *Assessment-Sensitivity: Relative Truth and its Applications*. Oxford University Press, Oxford, 2014.
- Aidan McGlynn. Believing things unknown. *Noûs*, 47(2):385–407, 2013.
- Aidan McGlynn. *Knowledge First?* Palgrave MacMillan, New York, 2014.
- Paul McNamara. Must I do what I ought (or will the least I can do do)? In Brown and Carmo, editors, *Deontic Logic, Agency, and Normative Systems*. Springer-Verlag, Berlin, 1996.
- Alfred Mele. Internalist moral cognitivism and listlessness. *Ethics*, 106:727–753, 1996.
- Alan Millar. What is it that cognitive abilities are abilities to do? *Acta Analytica*, 24(4): 223–236, 2009.
- Alexander Miller. *Introduction to Contemporary Metaethics*. Polity, Cambridge, 2003.
- Lisa Miracchi. Competence to know. *Philosophical Studies*, 172:29–56, 2015.
- Martin Montminy. Why assertion and practical reasoning must be governed by the same epistemic norm. *Pacific Philosophical Quarterly*, 94:57–68, 2013.
- G.E. Moore. Certainty. In *Philosophical Papers*. Allen and Unwin, London, 1959.
- G.E. Moore. *Commonplace Book: 1919-1953*. Allen and Unwin, London, 1962.
- Marcin Morzycki. Adjectival extremeness: Degree modification and contextually restricted scales. *Natural Language and Linguistic Theory*, 30(2):567–609, 2012.
- Sarah Moss. Scoring rules and epistemic compromise. *Mind*, 120(480):1053–1069, 2011.
- Sarah Moss. Time-slice epistemology and action under indeterminacy. In Hawthorne and Gendler, editors, *Oxford Studies in Epistemology*. Oxford University Press, Oxford, 2015a.

- Sarah Moss. On the semantics and pragmatics of epistemic vocabulary. *Semantics and Pragmatics*, 8(5):1–81, 2015b.
- Ram Neta and Guy Rohrbaugh. Luminosity and the safety of knowledge. *Pacific Philosophical Quarterly*, 85(4):396–406, 2004.
- Nikolaj Nottelman. The deontological conception of epistemic justification: A reassessment. *Synthese*, 190(12):2219–2241, 2013.
- Robert Nozick. *Philosophical Explanations*. Harvard University Press, Cambridge, 1981.
- Erik J. Olsson. A naturalistic approach to the generality problem. In Kornblith and McLaughlin, editors, *Goldman and his Critics*. Blackwell, Oxford, forthcoming.
- Robert Pasnau. Epistemology idealized. *Mind*, 122:987–1021, 2014a.
- Robert Pasnau. After certainty: A history of our epistemic ideals and illusions. Isaiah Berlin Lectures, 2014b.
- Carlotta Pavese. Know how and gradability, MS.
- Hazel Pearson. A judge-free semantics for predicates of personal taste. *Journal of Semantics*, 30:103–154, 2013.
- Philip Pettit and Michael Smith. Freedom in belief and desire. *Journal of Philosophy*, 93(9):429–449, 1996.
- Alvin Plantinga. *Warrant: The Current Debate*. Oxford University Press, Oxford, 1993.
- David Plunkett and Timothy Sundell. Disagreement and the semantics of normative and evaluative terms. *Philosophers' Imprint*, 13:1–37, 2013.
- Paul Portner. *Modality*. Oxford University Press, Oxford, 2009.
- Jesse Prinz. *The Emotional Construction of Morals*. Oxford University Press, New York, 2007.
- Duncan Pritchard. *Epistemic Luck*. Oxford University Press, New York, 2005.
- James Pryor. The skeptic and the dogmatist. *Noûs*, 34(4):517–549, 2000.
- Colin Radford. Knowledge: By examples. *Analysis*, 27:1–11, 1966.
- Baron Reed. A defense of stable invariantism. *Noûs*, 44(2):224–244, 2010.
- Baron Reed. Certainty. The Stanford Encyclopedia of Philosophy, winter 2011.
- Hans Reichenbach. *The Rise of Scientific Philosophy*. University of California Press, 1963.

- Steven Reynolds. Testimony, knowledge, and epistemic goals. *Philosophical Studies*, 110: 139–161, 2002.
- Michael Ridge. Ecumenical expressivism: Finessing Frege. *Ethics*, 116:302–336, 2006.
- Michael Ridge. Expressivism and epistemology: Epistemology for ecumenical expressivists. *Aristotelian Society Supplementary Volume*, 81(1):83–108, 2007.
- Gideon Rosen. Metaphysical dependence: Grounding and reduction. In *Modality: Metaphysics, Logic, and Epistemology*. Oxford University Press, Oxford, 2010.
- Jacob Ross and Mark Schroeder. Belief, credence, and pragmatic encroachment. *Philosophy and Phenomenological Research*, 88(2):259–288, 2014.
- Carmen Rotstein and Yoad Winter. Total adjectives vs. partial adjectives: Scale structure and higher order modifiers. *Natural Language Semantics*, 12:259–288, 2004.
- Bertrand Russell. *The Problems of Philosophy*. Henry Holt and Company, New York, 1912.
- Bertrand Russell. *Human Knowledge: Its Scope and its Limits*. George Allen & Unwin, London, 1948.
- Patrick Rysiew. The context-sensitivity of knowledge attributions. *Noûs*, 35(4):477–514, 2001.
- Jerry Sadock. On testing for conversational implicature. In P. Cole, editor, *Syntax and Semantics 9: Pragmatics*, pages 281–297. 1978.
- Galit Sassoon and Natalia Zevakhina. Granularity shifting: Experimental evidence from degree modifiers. *Proceedings of SALT*, 22:226–246, 2012.
- Uli Sauerland and Penka Stateva. Scalar vs. epistemic vagueness: Evidence from approximators. *Proceedings of SALT*, 17, 2007.
- Jonathan Schaffer. Knowledge in the image of assertion. *Philosophical Issues*, 18(1):1–19, 2008.
- Jonathan Schaffer. On what grounds what. In Chalmers and Wasserman, editors, *Metametaphysics*. Oxford University Press, Oxford, 2009.
- Frederick Schmitt. *Hume’s Epistemology in the Treatise*. Oxford University Press, New York, 2014.
- Mark Schroeder. *Ought*, agents, and actions. *Philosophical Review*, 120(1):1–43, 2011a.

- Mark Schroeder. What does it take to ‘have’ a reason? In Reisner and Steglich-Peterson, editors, *Reasons for Belief*. Cambridge University Press, New York, 2011b.
- Duns Scotus. *Philosophical Writings*. Hackett, Indianapolis, 1987.
- Anne-Marie Simon-Vandenberg and Karin Aijmer. *The Semantic Field of Modal Certainty: A Corpus-Based Study of English Adverbs*. Mouton de Gruyter, Berlin, 2007.
- Brian Skyrms. *Choice and Chance*. Wadsworth Publishing Co., Belmont, CA, 3rd edition, 1986.
- Aaron Sloman. ‘ought’ and ‘better’. *Mind*, 79:385–394, 1970.
- M.A. Slote. Assertion and belief. In Dancy, editor, *Papers on Language and Logic*. Keele University Library, Keele, 1979.
- Martin Smith. What else justification could be. *Noûs*, 44:10–31, 2010.
- Michael Smith. *The Moral Problem*. Blackwell, Oxford, 1994.
- Declan Smithies. The normative role of knowledge. *Noûs*, 46(2):265–288, 2012.
- Ernest Sosa. How to resolve the pyrrhonian dialectic: A lesson from descartes. *Philosophical Studies*, 85:229–249, 1997.
- Ernest Sosa. How to defeat opposition to moore. *Philosophical Perspectives*, 13:137–149, 1999.
- Ernest Sosa. Relevant alternatives, contextualism included. *Philosophical Studies*, 119: 35–65, 2004.
- Ernest Sosa. *Apt Belief and Reflective Knowledge*, volume 1. Oxford University Press, Oxford, 2007.
- Ernest Sosa. How competence matters in epistemology. *Philosophical Perspectives*, 24: 465–476, 2010.
- Ernest Sosa. *Judgment and Agency*. Oxford University Press, Oxford, 2015.
- Robert Stalnaker. *Context and Content*. Oxford University Press, New York, 1999.
- Jason Stanley. *Knowledge and Practical Interests*. Oxford University Press, New York, 2005a.
- Jason Stanley. Fallibilism and concessive knowledge attributions. *Analysis*, 65(2):126–131, 2005b.
- Jason Stanley. Knowledge and certainty. *Philosophical Issues*, 18(1):35–57, 2008.

- Will Starr. Dynamic expressivism about deontic modality. In Charlow and Chrisman, editors, *Deontic Modality*. Oxford University Press, New York, forthcoming.
- Tamina Stephenson. Judge dependence, epistemic modals, and predicates of personal taste. *Linguistics and Philosophy*, 30(4):487–525, 2007a.
- Tamina Stephenson. Indicative conditionals have relative truth conditions. *Proceedings of the Annual Meeting of the Chicago Linguistics Society*, 43(1):231–242, 2007b.
- Tamina Stephenson. *Towards a Theory of Subjective Meaning*. Ph.D. Thesis, MIT, 2007c.
- Matthias Steup. *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue*. Oxford University Press, Oxford, 2001.
- C.L. Stevenson. *Ethics and Language*. Yale University Press, New Haven, 1944.
- C.L. Stevenson. *Facts and Values*. Greenwood Press, Westport, CT, 1963.
- Michael Stocker. Desiring the bad: An essay in moral psychology. *The Journal of Philosophy*, 76:738–753, 1979.
- P.F. Strawson. Freedom and resentment. In *Freedom and Resentment and other Essays*, pages 1–25. Methuen, London, 1974.
- Timothy Sundell. Disagreement about taste. *Philosophical Studies*, 39(3):359–396, 2011.
- Jonathan Sutton. Stick to what you know. *Noûs*, 39(3):359–396, 2005.
- Jonathan Sutton. *Without Justification*. MIT Press, Cambridge, MA, 2007.
- Sigrun Svavarsdóttir. Moral cognitivism and motivation. *The Philosophical Review*, 108: 161–219, 1999.
- Eric Swanson. *Interactions with Context*. Diss., MIT, 2006.
- Kurt Sylvan and Ernest Sosa. The place of reasons in epistemology. In Star, editor, *The Oxford Handbook of Reasons and Normativity*. Oxford University Press, Oxford, 2014.
- Weng Hong Tang. Reliability theories of justified credence. *Mind*, forthcoming.
- Paul Tidman. Critical reflection: An alleged epistemic duty. *Analysis*, 56(4):268–276, 1996.
- Kelly Trogon. An introduction to grounding. In Hoeltje and Steinberg, editors, *Varieties of Dependence*. Philosophia Verlag, Munich, 2013.
- John Turri. Prompting challenges. *Analysis*, 70(3), 2010.
- John Turri. The express knowledge account of assertion. *Australasian Journal of Philosophy*, 89(1):37–45, 2011a.

- John Turri. Manifest failure: The gettier problem solved. *Philosophers' Imprint*, 11(8): 1–11, 2011b.
- Peter Unger. A defense of skepticism. *The Philosophical Review*, 80(2):198–219, 1971.
- Peter Unger. *Ignorance*. Clarendon Press, Oxford, 1975.
- Bas van Fraassen. *Laws and Symmetry*. Oxford University Press, New York, 1989.
- Robert van Rooij. Vagueness and linguistics. In Ronzitti, editor, *The Vagueness Handbook*, pages 123–170. Springer, Dordrecht, 2011.
- Mark van Roojen. Moral rationalism and rationalist amorism. *Ethics*, 120:495–525, 2010.
- Frank Veltman. Defaults in update semantics. *Journal of Philosophical Logic*, 25(3):221–261, 1996.
- Kai von Fintel and Thony Gillies. *Must... stay... strong!* *Natural Language Semantics*, 18(4):351–383, 2010.
- Kai von Fintel and Thony Gillies. Might made right. In Egan and Weatherson, editors, *Epistemic Modality*. Oxford University Press, Oxford, 2011.
- Kai von Fintel and Sabine Iatridou. What to do if you want to go to harlem: Anankastic conditionals and related matters, 2005. URL <http://web.mit.edu/fintel/www/harlem.pdf>.
- Kai von Fintel and Sabine Iatridou. How to say ‘ought’ in foreign: the composition of weak necessity modals. In Guéron and Lecarme, editors, *Time and Modality (Studies in Natural Language and Linguistic Theory 75)*. Springer, 2008.
- Brian Weatherson. Can we do without pragmatic encroachment? *Philosophical Perspectives*, 19:417–443, 2005.
- Brian Weatherson. $E \neq K$, 2009. URL <http://brian.weatherson.org/EK.pdf>.
- Brian Weatherson. Defending interest-relative invariantism. *Logos and Episteme*, 2:591–609, 2011.
- Brian Weatherson. Knowledge, bets, and interests. In Brown and Gerken, editors, *Knowledge Ascriptions*, pages 75–103. Oxford University Press, New York, 2012.
- Ralph Wedgwood. *The Nature of Normativity*. Oxford University Press, Oxford, 2007.
- Malte Willer. Dynamics of epistemic modality. *Philosophical Review*, 122(1):44–92, 2013.
- Timothy Williamson. Knowing and asserting. *Philosophical Review*, 105(4):489–523, 1996.

- Timothy Williamson. Knowledge as evidence. *Mind*, 106(424):1–25, 1997.
- Timothy Williamson. Conditionalizing on knowledge. *British Journal for the Philosophy of Science*, 49(1):89–121, 1998.
- Timothy Williamson. *Knowledge and its Limits*. Oxford University Press, Oxford, 2000a.
- Timothy Williamson. Scepticism and evidence. *Philosophy and Phenomenological Research*, 60(3):613–628, 2000b.
- Timothy Williamson. Replies to commentators. *Philosophical Books*, 45(4):313–323, 2004.
- Timothy Williamson. Contextualism, subject-sensitive invariantism and knowledge. *The Philosophical Quarterly*, 55(219):213–235, 2005.
- Timothy Williamson. The Amherst Lecture in Philosophy 4, 2009.
- Timothy Williamson. Knowledge first epistemology. In Bernecker and Pritchard, editors, *The Routledge Companion to Epistemology*. Routledge, New York, 2011.
- Timothy Williamson. Knowledge first. In Steup, Turri, and Sosa, editors, *Contemporary Debates in Epistemology*. Blackwell, Oxford, 2 edition, 2013.
- A.D. Woozley. Knowing and not knowing. *Proceedings of the Aristotelian Society*, 53: 151–172, 1952.
- Stephanie Wykstra. *A Defense of Cartesian Certainty*. Rutgers PhD dissertation, 2008.
- Seth Yalcin. Epistemic modals. *Mind*, 116(464):983–1026, 2007.
- Seth Yalcin. Probability operators. *Philosophy Compass*, 5(11):916–937, 2010.
- Seth Yalcin. Nonfactualism about epistemic modality. In Egan and Weatherson, editors, *Epistemic Modality*. Oxford University Press, Oxford, 2011.
- Seth Yalcin. Bayesian expressivism. *Proceedings of the Aristotelian Society*, 112(2):123–160, 2012a.
- Seth Yalcin. Context probabilism. *Proceedings of the 18th Amsterdam Colloquium*, pages 12–21, 2012b.
- Seth Yalcin. Epistemic modality *De Re*. *Ergo*, 2(19):475–527, 2015.
- Michael Zimmerman. Taking luck seriously. *Journal of Philosophy*, 99(11):555–576, 2002.