

Contents

Contents	iii
Introduction	3
1 'Knowledge'	11
1.1 Contextualism in General	11
1.2 Kaplan on Character and Content	12
1.3 Modals	14
1.4 'Knows' Contextualism and Skepticism	15
1.5 Elusive Knowledge	17
1.6 Quantifiers	18
1.7 Lewis and Lewisian Contextualism	19
1.8 Epistemic Standards	23
1.9 Invariantism with Shifting Standards	26
1.10 Factivity	29
1.11 Modality and Knowledge Ascriptions	30
1.12 Differences between Knowledge and Quantifiers	32
1.13 Knowledge Embedded in Conditionals	34
1.14 Is Contextualism Ad Hoc?	36
2 Sensitivity	38
2.1 Two Puzzles	39
2.2 David Lewis on Counterfactuals	41
2.3 Counterfactual Contextualism	43
2.4 Rules for Possibilities	45
2.5 Karen Lewis and Ignorance of Counterfactuals	47
2.6 Knowledge and Sensitivity	49
2.7 Equivocation and Necessary Conditions	52
2.8 Strengthening the Antecedent	54
2.9 Sensitivity, Safety, and Knowledge	54
3 Evidence	57
3.1 Motivation for E=K	57
3.2 An Argument Against E=K	59
3.3 Evidence as Non-Inferential?	60
3.4 Alexander Bird and 'Holmesian Inference'	61
3.5 Abominable Conjunctions	62
3.6 Non-Contextualist Responses	64

3.7	Contextualist E=K	66
3.8	The Intuitions Again	69
3.9	Evidence as Important	71
3.10	Circularity and Basic Knowledge	71
3.11	Lewis and Cartesian Contextualism	74
3.12	Moorean Contextualism	76
3.13	Skeptical Intuitions and Moorean Contextualism	79
3.14	Radical Skepticism	82
4	Justification	86
4.1	Initial Clarifications	87
4.2	Desiderata for a Theory of Justification	88
4.3	J=K?	91
4.4	Justification as Potential Knowledge	93
4.5	Is JPK Internalist?	94
4.6	Contextualism	96
4.7	Steven Reynolds	98
4.8	Alexander Bird	98
4.9	Justification as a Normative Status	100
4.10	An Objection	102
4.11	Lotteries	104
4.12	History	105
4.13	Reliability	106
	Appendix to Ch. 4	107
5	Action	113
5.1	Use of 'Knows'	113
5.2	Reasons	114
5.3	Contextualism and Norms	116
5.4	Intuitive Counterexamples to Necessity	118
5.5	Intuitive Counterexamples to Sufficiency	121
5.6	More Specific Theoretical Intuitions	125
5.7	The Thought-Bubble Model of Practical Reasoning	128
5.8	Counter-Closure	130
5.9	Locke on Ethical Theory	132
5.10	Schroeder on Ethical Theory	133
5.11	Reason-to	136
5.12	Internalism and External Redundancy	137
5.13	An Ethical Analogy	139
5.14	A Challenge to Internalist KR	140
5.15	Contextualism and Symmetry	143
5.16	Internalism and Basic Knowledge	147
6	Assertion	149
6.1	Stanley and the Certainty Norm	150
6.2	The Factivity Challenge	152
6.3	High-Standards Assertability of Low-Standards Knowledge	155
6.4	DeRose 2002	158

6.5	DeRose 2009	159
6.6	Contextualism and Norms, Again	160
6.7	The Method of Cases	162
6.8	KA and Good Enough Positions to Assert	164
6.9	Turri's 'Simple Test'	165
6.10	Incremental Assertion	167
6.11	Contexts and Possibilities	169
6.12	The Incremental Knowledge Norm of Assertion	170
6.13	Schaffer on Contrastivism and Assertion	174
6.14	Explaining Moore-Paradoxicality	176
7	Belief	178
7.1	Outright Belief	180
7.2	Shifty Data	183
7.3	Clarke, Sensitivism, and Belief as Credence One	184
7.4	Challenges for Clarke	187
7.5	Contextualism about Belief Ascriptions	189
7.6	'Knows' and 'Believes'	191
7.7	Knowledge and Proper Belief	197
7.8	Doxastic States and Epistemology	198
	Bibliography	200

Introduction

I begin with an analogy.

Diary of a Narcissist

Here is a diary entry by Reginald, a confused narcissist.

Dear Diary,

I am perturbed. As you know, I've long found myself, if not perhaps perfection, surely the next best thing to it; I thank Providence every day for raising me up so far above the vulgar. It is no exaggeration to say that hitherto, I have counted myself among the very most beautiful and significant people in all of Creation. But today I happened across a paper by a philosopher called David Kaplan. What I found there shook my deepest convictions to the core. For Kaplan argued that certain words—'demonstratives' or 'indexicals', he calls them—are *context sensitive*; the referents of these terms can *vary* according to the conversational context in which they're used. My first instinct was a sanguine one; his seemed an interesting and plausible semantic thesis. The referent of the word 'that', for example, is simply whatever it is at which my flawless finger happens to be pointing at a given moment when I'm speaking.

But it's not just that! It's one thing to recognise the general semantic framework—it's quite another to make particular entries in the list of context-dependent terms. Among Kaplan's list of context-dependent terms are the very dearest and most important to me! He includes on his list, for example, such touchstones as 'I' and 'me'! Can you imagine, diary? *I*—Reginald the all-right—dependent on such contingencies as conversational contexts? Never in my wildest dreams would I have imagined that anyone would so trivialise *me*. Needless to say, I am deeply shaken. Must I accept that I am so unimportant? That there is nothing special about me, but rather than I'm merely, *whoever* happens to be speaking in a given conversation? The thought terrifies me. Tomorrow I shall attempt to rebut Kaplan's defamatory arguments; tonight I am too shocked. I must rest.

Fondly,

Reginald

Reginald's error is not difficult to diagnose; he's very bad at the use-mention distinction. At times, in the passage above, he is *using* words like 'I' and 'me', thus talking about himself; at other times he's *mentioning* them, thus talking about those English *words*. (You are a person, not an English word; 'you' is an English word, not a person.) Kaplan (1977) gives a context-sensitive theory about the *word* 'I'; but in doing so, he doesn't give a theory about me or Reginald or anybody else. Saying

that ‘I’ just picks out whoever happens to be speaking isn’t tantamount to saying that anybody is unimportant. Kaplan’s theory of demonstratives does not imply that Reginald isn’t special. So much is, I take it, pretty obvious.

Nevertheless, when epistemologists start thinking about contextualism about ‘knows’—roughly, the thesis that ‘knows’ is similar to words like ‘I’ and ‘me’ in that its referent depends on the conversational context—the corresponding point is not always treated as quite so obvious. For example, echoing Reginald, I have often encountered a perceived tension between contextualism about ‘knows’ and the idea that knowledge is important. I encounter this perception more often in conversation than in print, but Alvin Goldman does give a brief expression to a version of it here:

A popular view in contemporary epistemology (with which I have much sympathy) is that knowledge has an important context-sensitive dimension. The exact standard for knowledge varies from context to context. Since it seems unlikely that natural kinds have contextually variable dimensions, this renders it dubious that any natural kind corresponds to one of our ordinary concepts of knowledge. Goldman (2007, p. 17)¹

Assuming that the ‘popular view’ in question is contextualism, Goldman’s fallacy is the same as Reginald’s: it is a use–mention error. There is no straightforward connection between the semantic properties of the English word ‘knows’ and the metaphysical properties of knowledge. (Compare the fact that there is no straightforward connection between Kaplan’s observations about indexicals and Reginald’s beliefs about himself.)²

This book is about the relationship between contextualism about ‘knows’, on the one hand, and epistemological theorising about knowledge, on the other. It is a mistake to think that there is any very straightforward connection between them, but is there a subtler one? I shall suggest that there is. In particular, I will argue that there is a mutually supporting package of views, combining a particular brand of contextualism about ‘knows’ with a particular interpretation of the ‘knowledge first’ programme, according to which knowledge is a theoretically fundamental and important mental state. Assuming contextualism, the sentence, ‘knowledge is a theoretically fundamental and important mental state’ may be a context-sensitive one—this sentence could be used to express different propositions in different conversational contexts. That doesn’t mean it isn’t true and informative and theoretically enlightening, or that we can only mention, rather than use it. But it does mean that when we use it, we must use it carefully, and attend to potential ambiguities.

I have written a book about the relationship between contextualism about ‘knows’ (hereafter ‘contextualism’) and the knowledge first programme because I think both of these views have much to commend them. I have defended versions of both views in print—many of the ideas from these earlier papers are incorporated into this book. Along the way, I will say something about why I find these two views attractive, and I hope that many readers will come to look sympathetically on them, but it is not my primary purpose to argue for either contextualism or the knowledge first stance. Rather, I hope to show that these disparate views, though independently developed and considered, and widely thought to stand in a kind of tension, in fact fit rather well together. Contextualism can help the knowledge-first theorist respond to certain important objections to that stance; knowledge-first can help provide a theoretical motivation for the contextualist’s claims. Any philosopher who wants to consider one of these views would do well to do so along with the other. Throughout

¹In work in progress, Dani Rabinowitz also defends a version of Goldman’s argued incompatibility.

²Compare also the remarks of Jenkins (2008, p. 70): ‘Note that none of this context sensitivity [about ‘explanation’ and related terms] gives us any reason to be mind-dependent anti-realists about explanation. What depends on contexts (and hence on the intentions, interests and so on of the utterer and/or audience) is what is expressed by terms like ‘good explanation’ and ‘best explanation’, *not* what counts as a good or bad explanation *once these aspects of meaning are settled*. The thought is also related to Keith DeRose’s ‘intellectualist’ motivation for contextualism—see DeRose (2009, ch. 6). However, I will suggest in Ch. 1 that this motivation is not mandatory for contextualists.

the book I'll canvass a series of studies of particular issues, exploring particular ways in which contextualism and the knowledge first stance may synergise. I include a chapter summary at the end of this introduction. First, however, I should set out the important backdrop that is the knowledge first program.

Knowledge First

It's not particularly controversial that knowledge is epistemologically interesting; but in what way is it interesting? A prevalent assumption in some of the history of philosophy had it that knowledge is a central *explanandum* in epistemology: a central task is to explain what knowledge is, how or whether knowledge can be attained, whether and why knowledge is particularly valuable, etc. Paradigmatic of this approach is the "theory of knowledge" literature spawned by Edmund Gettier's famous paper. The aim of this literature was to provide an "analysis" of knowledge in more fundamental terms—to explain knowledge, for example, in terms of belief, justification, truth, evidence, etc.

More recently, some philosophers have attempted to approach epistemological questions concerning knowledge from a different angle. Timothy Williamson is widely credited with this change in perspective; his 2000 book, *Knowledge and Its Limits*, advocated for an approach he calls "knowledge first".³ The knowledge first stance reverses the traditional order of explanation: knowledge is treated as *explanans*, rather than as explanandum. The idea here is that knowledge is in some sense fundamental, and can be used to illuminate other states of epistemological interest. For example, perhaps we should understand evidence in terms of knowledge, instead of vice versa.

In our (2016), Carrie Jenkins and I distinguished between a number of distinct knowledge-first theses. In particular, we argued for a clear distinction between the *metaphysical* and *representational* claims that travel under that banner. Metaphysical knowledge first claims are claims about knowledge itself; questions about whether knowledge is a mental state, or whether it is a (relatively or absolutely) fundamental feature of reality are foregrounded. Representational claims, by contrast, have to do with how we *think* about or *talk* about knowledge; questions about whether the concept KNOWS has BELIEVES as a component, or whether knowledge ascriptions typically or invariably proceed by virtue of belief ascriptions, for instance, characterise the issues in discussing representational knowledge first theses.

Although Williamson defends views of both families, and although some authors have argued for strong connections between them, Jenkins and I argued that the views are prima facie independent. The version of the 'knowledge first' programme I am interested in exploring in this book belongs to the *metaphysical* family. I am more interested, for example, in the idea that knowledge is a mental state that has theoretically significant roles in explaining things like action, belief, and justification, than I am in the idea that knows is a mental state concept, or the question whether knowledge ascriptions proceed via tacit belief ascriptions. Throughout the book, I will be interested in exploring ways in which knowledge connects to other areas. My sympathies lie with the broad methodology of Lewis (1984), according to which appearance in good theorising about the world is a mark of fundamentality.

Lewis himself focused on the *perfectly* natural; his view is that the perfectly natural properties are the ones that appear in the basic elements of a minimally adequate theory of the world. In my view, it is reasonably natural to extend that thought to the idea that *relatively* fundamental properties appear in less fundamental, but genuinely real, theories.⁴ The idea of the knowledge first project, so

³Contemporary pre-cursors included Zagzebski (1994), Zagzebski (1996), and Craig (1990). There were certainly earlier precedents, e.g. Kneale (1949, p. 18): 'According to the view presented here, knowledge is sui generis and the two varieties of belief are to be defined by reference to it.' For an overview and discussion of historical knowledge-first ideas, see Marion (2000) and Mulligan (2014).

⁴The same, I think, goes for Ted Sider's more recent treatment of similar themes—see Sider (2011, p. 29).

interpreted, is that knowledge is *more metaphysically fundamental* than, for example, was supposed in the project of trying to explain knowledge in terms of justification and belief. This metaphysical claim should be understood in a way similar to the way that many of us think that green is metaphysically prior to grue, or that electrons are metaphysically prior to laptop computers. The approach is plausible to the degree to which metaphysical theorising in terms of knowledge is fruitful. This is why, for example, Timothy Williamson's suggestion that citing knowledge typically makes for a better explanation than does citing belief, in the explanation for someone's action⁵; that it appears in the best theory of action explanation is some reason to think knowledge is of some theoretical significance, hence metaphysical fundamentality.⁶ So likewise for other theoretical roles posited for knowledge—perhaps one need invoke knowledge in order to explain evidence, or justification, or mental content, etc.

If a picture along these lines is right, then finding such roles for knowledge will constitute evidence that knowledge is a relatively fundamental state. To be sure, the Lewisian picture about fundamentality is controversial; adjudicating the serious metaphysical question about the relationship between theoretical roles and fundamentality is well beyond my present scope.⁷ Even setting such relations aside, the idea that knowledge connects in deep theoretical ways with action, justification, belief, evidence, etc. is itself of significant interest.

If contextualism is correct, then the sentence, 'knowledge is a relatively fundamental state' may itself be a context-sensitive sentence. (I'll explain contextualism in much more detail in Ch. 1.) Contextualists need to have something to say about how it is intended. As will emerge throughout the book, I think that different proposed theoretical roles for knowledge ought to be treated differently in this respect. There is a kind of general perception among epistemologists that a knowledge first stance fits poorly with contextualism. I think this is a mistake.

Contextualism 'Evading' Epistemology?

One preliminary worry has to do with the general relevance of contextualism to epistemology. In effect, it starts with the observations I made above about the independence of claims about 'knows' ascriptions from claims about knowledge, and concludes from this that insofar as we care about the latter, we ought to ignore the former. As Ernest Sosa puts it, '[c]ontextualism replaces a given question [about knowledge] with a related but different question [about "knows" ascriptions]'. Sosa (2000, p. 1) And it is not at all clear that an answer to the latter question will bear on the former. To adapt one of Sosa's own pithy examples, Patience might say to herself:

I am very confident that people often utter truths when they say "Somebody loves me."
But does anybody at all love me?⁸

In exactly the same way, one may be convinced by the truth of contextualism, and rest easy that people often utter truths when making knowledge ascriptions⁹; still, one might wonder whether anybody at all knows anything. (If you're having a hard time seeing how this could be an open question, let's stipulate that I'm writing in at least a moderately skeptical context now; the thought is that people utter truths with such ascriptions only in more lax contexts.) It is important to keep clear the distinction between *object language* discussion, that *uses* terms to talk about the world

⁵Williamson (2000, pp. 62–3)

⁶See Ichikawa and Jenkins (2016, §§2.1, 4) for further discussion.

⁷For some discussion, see Jenkins (2013).

⁸Sosa (2000, p. 3)

⁹As I use the term, knowledge ascriptions are sentences using the word 'knows'.

more broadly, and the *meta-language*, that *mentions* terms in order to discuss them.¹⁰ So it is that Sosa raises the challenge of how it is that contextualism can be relevant for epistemology. Hilary Kornblith has gone so far as to argue that contextualism is ‘largely irrelevant to epistemological concerns’.¹¹

Contextualism has two kinds of responses available to this worry. One of them is the thesis of this book: that a holistic theory combining contextualism and the uncontroversially epistemologically-relevant knowledge first programme is appealing; if contextualism helps make that programme more plausible, it is contributing helpfully to epistemology.

But there is also a more schematic reply available, which I’d like to articulate now. I agree with Sosa’s claim that contextualism doesn’t bear in any direct way on standard epistemological questions about the nature and extent of human knowledge. But I dispute the inference from this point to its *irrelevance* for the latter. For even though it doesn’t provide any straightforward evidential support for any particular epistemological view, contextualism is, if true, crucially important for the *methodology* of epistemology. Anyone interested in understanding knowledge has an interest in thinking clearly about knowledge, and if contextualism is true, then equivocation on ‘knows’ is possible. So if contextualism is true, epistemologists must exhibit sensitivity to this fact.¹²

Take for example a classical skeptical argument like this one: I have no way to tell whether or not I will unexpectedly drop dead tonight; therefore I don’t know much at all about what I may or may not do tomorrow. Its premise enjoys some intuitive plausibility, but its conclusion is far more skeptical than most epistemologists want to admit. It is tempting to suppose that one has to choose between these attractive ideas; then non-skeptics are burdened with the task of explaining away the attractiveness of the initial claim of ignorance. But if contextualism is right, then one needn’t reject either intuitive starting-point; the argument to the effect that one must is equivocal. Careful attention to the language we use is sometimes the only way to avoid confusion about that which our language is about.¹³

This is in effect the same observation that Timothy Williamson makes in a different context in *The Philosophy of Philosophy*:

Philosophers who refuse to bother about semantics, on the grounds that they want to study the non-linguistic world, not our talk about that world, resemble scientists who refuse to bother about the theory of their instruments, on the grounds that they want to study the world, not our observation of it. Such an attitude . . . produces crude errors. Williamson (2007, pp. 284–85)

I do not wish to deny that there is also an important sense in which contextualism is not *the answer* to deep epistemological questions—I’ll discuss this in much more depth in Ch. 3. But there is every reason to expect it to be *of relevance*.¹⁴

¹⁰The meta-language uses *some* terms—consider for example this sentence: “‘Me’ is an example of a context-sensitive term.” This sentence *mentions*, rather than *uses* the word ‘me’, which is why it is about that word instead of me (i.e., Jonathan). But it also *uses* the word ‘example’ and six other words.

¹¹Kornblith (2000, p. 24). As an OUP referee points out, in addition to the critique given in the main text, it seems that Kornblith’s argument depends on the dubious (to understate things!) assumption that consideration of skepticism exhausts epistemology’s concerns.

¹²See DeRose (2009, pp. 18–19) for a version of this point.

¹³Indeed, Sosa himself recognises as much in different contexts. Compare his remarks in Sosa (2009, p. 104) about a different kind of semantic issue: ‘Semantic ascent does have a place in epistemology if only when we attempt to understand persistent disagreement by appeal to ambiguity or context-dependence.’

¹⁴Sosa also gives another argument against the epistemic significance of contextualism in his paper—one that relies on the assumption that contexts in which one engages in epistemology are inevitably skeptical ones relative to which all or nearly all knowledge ascriptions are false. I dispute this assumption in §6.2. See also Blome-Tillmann (2014, pp. 49–52).

Context-Sensitive Normative Discourse

A second preliminary objection to my project is more specific to the interaction between contextualism and knowledge norms. A significant component of the knowledge first stance relates knowledge to normative concepts. For example, Williamson (2000) argues that knowledge is the constitutive epistemic norm of assertion, which implies something like this:

(N) If and only if S knows p , S is epistemically permitted to assert that p

Other knowledge norms have also been proposed, many of which I will consider in detail in this book. But all of them, according to the objection I am now considering, fit badly with contextualism. The objection runs like this: any contextualist who adopts a knowledge norm is committed to problematic contextualism about normative concepts. Suppose, for example, that a contextualist endorses (N). Since the left-hand side is context-sensitive, by the contextualist's lights, our contextualist must also be a contextualist about the right-hand side too. But (the objection continues) it is not plausible to endorse this kind of normative contextualism in the cases at issue.¹⁵

The argument is fallacious; it semantically ascends and descends freely, and a contextualist need not accept it. Semantic ascent is the move from an object-language claim to a metalanguage claim, as in the move from the claim that somebody loves me to the claim that an utterance of 'somebody loves me' is true; semantic descent is a matter of *disquotation*, inferring an object-language claim from a metalanguage claim: the utterance of 'somebody loves me' is true, therefore somebody loves me. Neither move is generally valid for context-sensitive discourse. (Suppose that nobody loves me, but that some people love Patience. And suppose further that Patience is the person who says 'somebody loves me'. Then that sentence is true as uttered, even though nobody loves me.)

Returning to the case of (N), observe first that it is stated in the object language; it *uses* words like 'knows' and 'permitted'; it does not mention them. It is not *about* these words. (There are no quotation marks in (N).) Therefore the move from (N) and contextualism about the left-hand side to contextualism about the right-hand side is invalid. Only a strong metalinguistic generalisation of (N) could have these kinds of consequences.

In fact, there are at least five options for contextualists who accept object-language claims like (N):

1. Decline to endorse any metalinguistic principle, instead interpreting (N) as holding in a particular favored context.
2. Interpret the object-level claim as applying in S's context, as does DeRose (2009) about assertion.
3. Endorse the principle in full metalinguistic generality, positing context-sensitivity in the normative language, as per the objection suggested.
4. Endorse the principle in full metalinguistic generality, positing context-sensitivity in the relevant 'assert' language.
5. Endorse the principle in full metalinguistic generality, holding that the context-invariant right-hand-side applies any time the left-hand-side is true in any context.

¹⁵Williamson (2005b) and Hawthorne (2004, pp. 86–9) each give arguments in this neighbourhood against contextualism. While both these authors have subtle things to say about the relationship between contextualism and knowledge norms, I think both underestimate the extent of the contextualist's resources.

These represent a catalogue of available tools to the contextualist; there is no reason a contextualist need adopt a uniform treatment for all knowledge norms or other theoretical principles involving knowledge. I will apply diverse strategies through the book, applying a version of (2) to the knowledge norm of assertion, a version of (5) to certain proposed connections between knowledge and rational action, and a version of (4) to the equation of knowledge and evidence, and to the relationships between knowledge, justification, and belief.

The point of this section was to give one flavour of the kind of tension that some theorists have perceived between contextualism and knowledge-first epistemology. Several similar tensions (and some less similar ones) will be explored throughout the book.

Outline of the Book

Although contextualism and the knowledge first project have developed independently, and are typically thought to stand in tension with one another, I will argue that this perceived tension is illusory. On the contrary, I mean to make the case in this book for the idea that contextualism and the knowledge first project are complementary: each has something to offer the other. The details will come over the course of the book, but one common theme will be that contextualism helps the knowledge first project to avoid counterintuitive consequences, while the knowledge first project can help certain forms of contextualism answer the challenge of being ad hoc.

I begin in Chapter 1 with a development of the contextualist semantics I prefer. It is inspired by and related to David Lewis's relevant alternatives approach to knowledge ascriptions—it descends from an approach I first developed in Ichikawa (2011a). According to the Lewisian idea, the context-sensitivity of knowledge ascriptions is modelled on the context-sensitivity of modals or sentences involving quantifiers—satisfying 'knows p ' requires evidence that conclusively rules out 'all' not- p cases, but where the domain of the 'all' depends in part on the conversational context. I'll argue in that chapter that a version of this idea captures quite a lot of the intuitive data, and avoids some of the challenges that have been levelled against contextualists. I'll also argue that there is a significant class of under-explored linguistic data, concerning the interaction of knowledge ascriptions with conditionals, which further motivates this kind of contextualist approach.

Chapter 2 picks up on the idea that knowledge requires a certain kind of counterfactual connection to the truth. This rather natural idea played significant roles in twentieth-century theorising about knowledge, but is almost entirely discredited today—not least because it is thought to countenance 'abominable conjunctions'. However, the contextualist approach developed in Ch. 1, when combined with an independently-motivated contextualist approach to counterfactual conditionals, allows the proposed connection to be seen in a new light. I'll argue that sensitivity, suitably understood, *is* a genuine necessary condition for knowledge; given contextualism, counterintuitive consequences—including abominable conjunctions—can be avoided.

In Chapter 3, I turn to two pressing questions about evidence—one concerns Timothy Williamson's famous suggestion that a subject's evidence comprises all and only her knowledge ('E=K'). I agree with the letter of this equation (suitably embedded within contextualism); but I also suggest that there is good reason to recognise a privileged category of *basic* knowledge/evidence. This motivates the second central question of the chapter: what is the best way to understand basic evidence? The contextualist approach described in Ch. 1 had it that the truth of a knowledge ascription requires *evidence* that rules out all relevant counter-possibilities—but this could be understood in terms of various approaches to evidence itself. One of the under-recognised commitments of David Lewis's contextualism—and of quite a lot of contemporary philosophy—is a kind of 'Cartesian' approach to basic evidence, whereby a subject's subjective experiences are the things that are known first and best. I'll challenge this assumption in this chapter, arguing that it's certainly not mandatory, and

very plausibly ill-motivated. I'll consider instead a kind of neo-Moorean *disjunctivist* approach to evidence, according to which that which a subject can see for herself—including propositions about the external world—are basic evidence playing a foundational role. Consequently, the relationship between contextualism and radical skepticism ends up on my view rather different than that which most contextualists so far have posited.

Chapter 4 develops a knowledge-first theory of epistemic justification within the contextualist framework. Its starting point is the approach to justification defended in Ichikawa (2014), according to which justification is a matter of 'potential knowledge'—I characterised this notion as a matter of being intrinsically identical to a possible subject with knowledge. Ch. 4 generalises from that approach in two ways: first, it motivates and explicates a metasemantic generalisation of that view. The result is a contextualist semantics for ascriptions of epistemic justification to match that given for knowledge. Second, it relaxes the internalist assumption I'd previously made, considering various ways to understand the notion of 'potential knowledge', corresponding to various conceptions of basic evidence. One upshot of this generalisation is a new theoretical understanding of the kind of approach that could even motivate *radically* externalist theories of justification, such as the identification of justification with knowledge.

The final three chapters turn to proposed knowledge norms. Chapter 5 considers and defends the knowledge norm of practical reasoning, and suggests that, contrary to many authors' claims, it does not have radical externalist implications about rational action; relatedly, it also doesn't, when combined with contextualism, require any sort of implausible contextualism about normative discourse about action. I'll suggest that arguments to the contrary often tacitly assume a highly questionable approach to practical reasoning, which I call the 'thought bubble model'. By contrast, in Chapter 6, I *will* defend a kind of subject-relative version of the knowledge norm of assertion, justifying and systematising the connection by reference to Stalnakerian models for conversational contexts. Finally, Chapter 7 defends a contextualist semantics for belief ascriptions, and argues for a systematic normative relation between belief and knowledge, and between 'believes' and 'knows'.