


*Blackwell
Companions to
Philosophy*

A COMPANION TO DAVID LEWIS



Edited by
BARRY LOEWER AND
JONATHAN SCHAFFER

WILEY Blackwell

Lewis on Knowledge Ascriptions

JONATHAN SCHAFFER

David Lewis's primary contribution to the theory of knowledge is his account of knowledge ascription, which integrates an elegant version of relevant alternatives theory with a detailed version of contextualism. His account is prefigured in his discussion of accommodation in "Scorekeeping in a Language Game" (1979), and forms the central topic of his "Elusive Knowledge" (1996). I will review his account (section 30.1), and then discuss some epistemic (section 30.2) and semantic (section 30.3) questions arising. This will be a friendly critique: I am sympathetic with the outlines of Lewis's contextualist relevant alternatives theory, but would amend some details.

Before proceeding, I should mention that Lewis offers discussions bearing on knowledge in various places, including:

- the treatment of conventions as common knowledge, provided in *Convention* (1969);
- the semantics of 'whether'-clauses, including their behavior in 'know whether' constructions, found in "Whether Report" (1982);
- the concern that modal realism entails modal skepticism, discussed in *On the Plurality of Worlds* (1986);
- the claim that knowing what it is like to have an experience does not consist in having information but only in having abilities, defended in "What Experience Teaches" (1988a); and
- the argument that we cannot know which properties occupy the theoretical roles of our best theory, presented in "Ramseyan Humility" (2009).

In these discussions Lewis tends to use the notion of knowledge in an intuitive way, and it is the burden of his contextualist relevant alternatives theory to fit this intuitive usage. I leave open the exegetical question of whether Lewis's theory always fits his usage.¹

I should also mention that Lewis offers various extended discussions of "formal epistemology," including:

- the treatment of objective chance via the Principal Principle, articulated in "Subjectivist's Guide to Objective Chance" (1980);

- the argument that treating desire as belief conflicts with standard decision theory, developed in “Desire as Belief” (1988b);
- the diachronic Dutch book argument for updating by Bayesian conditionalization, motivated in “Why Conditionalize?” (1999); and
- the argument that Sleeping Beauty ought to assign a credence of .5 to heads when she wakes, given in “Sleeping Beauty: Reply to Elga” (2001).

I leave open what if any connections there might be between Lewis’s treatment of formal epistemology and his treatment of “traditional” matters of knowledge and its ascription, save to note that he (1996, 553) characterizes our concept of knowledge as a “handy but humble approximation” which only provides “a very sloppy way of conveying very incomplete information about the elimination of possibilities.”

One might think that Lewis, by speaking of our concept of knowledge as a “handy but humble approximation,” means to marginalize the topic. But he (1996, 553) immediately adds that the concept may yet be “indispensable in practice,” and he evidently deems it worthy of detailed discussion. Indeed a great deal of his overall philosophical energy went into developing philosophically rigorous accounts of shifty but handy concepts such as counterfactuals and causation.² So I think that he regards knowledge as one of the many shifty but handy concepts for which a rigorous philosophical treatment is a difficult but necessary task.

30.1 Lewisian Knowledge Ascription

Lewis offers an account of knowledge ascription, which integrates relevant alternatives theory with contextualism. I begin with a review. This will involve presenting both his elegant version of relevant alternatives theory (section 30.1.1), and his detailed version of contextualism (section 30.1.2), and then displaying the combined account and its claimed virtues (section 30.1.3).

30.1.1 Lewisian Relevant Alternatives Theory

Lewis’s account begs to be misunderstood. From his title (“Elusive Knowledge”) up to the very end of his paper – including the “official statement” of his view – Lewis speaks *as if* he were presenting a theory of knowledge. Only at the end does he (1996, 566–7) mention that he has offered a simplified and potentially misleading presentation. But he still offers no corrected statement, instead trusting the reader who senses something amiss to appreciate how his account could have been stated in a more cumbersome but correct form, via semantic ascent. This is a lot of trust to place in the reader. So – if only to help the reader repay this trust – it will prove useful to state Lewis’s account correctly.

Thus the first point to clarify is that Lewis (despite simplified statements as if otherwise) is not really presenting an account of *knowledge*, but rather of *knowledge ascription*. He is not really saying when *s* knows that *p*, but rather when a sentence of the form ‘*s* knows that *p*’ is true relative to a context *c*. These are connected but distinct topics.³ As such he is *not* really saying that knowledge is elusive, or is destroyed by doing epistemology, or anything like that (despite seeming to say just these things). Properly understood, he is really saying something metalinguistic, about the contexts in which knowledge ascriptions come out true.

With this simplification firmly in mind, one can approach Lewis’s (1996, 551) official statement: “Subject *S* knows proposition *P* iff *P* holds in every possibility left uneliminated by *S*’s evidence; equivalently, if *S*’s evidence eliminates every possibility in which not-*P*.” The context dependence is

said to arise from the context dependence of the domain associated with ‘every’, which he (1996, 554) incorporates as: “*S* knows that *P* iff *S*’s evidence eliminates every possibility in which not-*P* – Pstt! – except for those possibilities that we are properly ignoring.” Semantically ascending, one reaches the more cumbersome metalinguistic claim that correctly states Lewis’s account:

Lewisian relevant alternatives theory: A sentence of the form ‘*s* knows that *p*’ is true in context *c* iff *s*’s evidence eliminates every not-*p* possibility relevant in *c*.

Lewisian relevant alternatives theory is distinctive in at least two main respects, the first of which – though obscured in Lewis’s simplified presentation – is its metalinguistic character.⁴ This proves crucial for integration with contextualism. With an object language account, contextualism would lead straight to contradiction, since one could derive both *s* knows that *p* (from a “favorable” context) and then its negation (from an “unfavorable” context). A metalinguistic account only allows one to derive that a given sentence (of the form ‘*s* knows that *p*’) is true in one context and false in another, and that is not only consistent and unproblematic, but just what one expects from contextually sensitive discourse. (By way of comparison, it is unproblematic and indeed expected that – due to the context sensitivity of ‘today’ – the sentence ‘today is Friday’ is true in a Friday context but false in a Monday context.)

The metalinguistic character of *Lewisian relevant alternatives theory* allows for contextualism but does not demand it. Context invariance is equally allowed insofar as nothing yet has been said about the determinants of relevance (Lewis’s account of relevance will be discussed in section 30.1.2), and so nothing yet rules out the prospect that exactly the same relevant alternatives might be generated in every single context, given a fixed subject *s* and proposition *p*. Thus it is strictly compatible with *Lewisian relevant alternatives theory* to hold that, for any subject *s* and proposition *p*, the relevant alternatives in every context are (i) every not-*p* possibility which is unrestrictedly logically possible (Unger 1975), or (ii) every not-*p* possibility with nonzero objective chance in *s*’s situation (Dretske 1991, 192), or even (iii) every not-*p* possibility which is practically relevant to *s* (Hawthorne 2004, 158–61; Stanley 2005, 85). Thus *Lewisian relevant alternatives theory* is usefully neutral on contextualism.

A second main respect in which *Lewisian relevant alternatives theory* is distinctive is in its elegance, in making the elimination of the relevant alternatives the *one and only* requirement. No other requirements – such as justification, truth, belief, or basing – are imposed.⁵ Instead Lewis uses his account of relevance – to which I now turn – to simulate some of these requirements, in distinctive ways.

30.1.2 *Lewisian Contextualism*

Lewis combines the elegance of *Lewisian relevant alternatives theory* with a detailed contextualist account of the determinants of relevance. Indeed a large portion (pp. 554–61) of “Elusive Knowledge” consists in discussing rules for determining the relevance of a possibility at a context. Lewis offers seven such rules:

Actuality: the possibility that actually obtains for the subject is always relevant;

Belief: any possibility that the subject believes or ought to believe obtains is always relevant;

Resemblance: any possibility that saliently resembles a relevant possibility (made relevant by any rule other than *Resemblance* itself) is always relevant;

Reliability: possibilities concerning errors in reliable processes (such as perception, memory, and testimony) are defeasibly irrelevant;

Method: possibilities concerning errors in sampling and in abduction are defeasibly irrelevant;

Conservatism: possibilities that are conventionally ignored are defeasibly irrelevant; and *Attention*: any possibility explicitly under discussion is always relevant.

Putting these seven rules together yields:

Lewisian contextualism: The relevance of a possibility at a context is determined by all and only the rules of *Actuality*, *Belief*, *Resemblance*, *Reliability*, *Method*, *Conservatism*, and *Attention*.

Lewis's rules can be usefully divided in several ways, of which I will mention three. First, there are factors that push towards expanding the relevant possibilities (*Actuality*, *Belief*, *Resemblance*, *Attention*), and factors that push towards contracting them (*Reliability*, *Method*, *Conservatism*). Second, there are factors that push with irresistible force (*Actuality*, *Belief*, *Resemblance*, *Attention*), and those that push with resistible force possibly overcome by a push in the other direction (*Reliability*, *Method*, *Conservatism*).⁶ Though – as will emerge in section 30.3 – Lewis largely undoes this second distinction at the end of his paper, allowing that all the factors he cites are resistible. Third – and most importantly for present purposes – there are invariant factors that are fixed across attributors (*Actuality*, *Belief*, *Reliability*, *Method*, *Conservatism*), and contextual factors that can vary across attributors (*Resemblance*, *Attention*). It is the postulation of these latter factors that make Lewis's account contextualist. Lewis's rules minus *Resemblance* and *Attention* do not allow for any two contexts to differ in their relevant possibilities, given a fixed subject *s* and a fixed proposition *p*.

Lewisian contextualism is distinctive in at least two main respects, the first of which is its attention to detail. Previous relevant alternatives theorists – whether contextualists or not – had left the notion of relevance at a largely intuitive level, leading previous commentators such as Sosa (1986, 585) to speak of relevance as “unacceptably occult.” While *Lewisian contextualism* leaves some details open,⁷ it achieves a level of precision that might be thought appropriate to the topic, and at any rate clearly surpasses anything found before.

A second main respect in which *Lewisian contextualism* is distinctive is its treatment of factors like justification, truth, and belief (this is the counterpart of the elegance of *Lewisian relevant alternative theory*, which did without mention of these matters). The treatment of truth turns out to be distinctive only in implementation, not in substance. Lewis imposes the usual requirement that the proposition *p* be true, just indirectly via the rules of relevance instead of directly via a condition on knowledge ascriptions.⁸ But the treatments of justification and belief are distinctive in substantive ways. For Lewis (1996, 551) denies that justification is necessary for knowledge, and he (1996, 556) likewise denies that belief is necessary for knowledge.⁹ Instead his *Belief* rule is supposed to function as an improved replacement for these requirements, giving belief and justification a minor role in generating relevance (as Lewis (1996, 556) notes: “This is the only place where belief and justification enter my story”), while still allowing for a true knowledge ascription in the absence of either belief or justification, given sufficient evidence for ruling out whatever alternatives got generated.

30.1.3 Lewisian Knowledge Ascription

Lewis's full theory of knowledge ascription is then the result of conjoining *Lewisian relevant alternatives theory* with *Lewisian contextualism*:

Lewisian knowledge ascription: A sentence of the form ‘*s* knows that *p*’ is true in context *c* iff *s*'s evidence eliminates every not-*P* possibility relevant in *c*, where the relevance of a possibility in *c* is determined by all and only the rules of *Actuality*, *Belief*, *Resemblance*, *Reliability*, *Method*, *Conservatism*, and *Attention*.

To illustrate, consider the following sentence:

- (1) Ann knows that there is a goldfinch in the garden.

By *Lewisian knowledge ascription*, (1) is true in a given context c iff Ann's evidence eliminates every possibility relevant in c in which it is not the case that there is a goldfinch in the garden. For instance, suppose that in context c_{easy} the only relevant possibility in which it is not the case that there is a goldfinch in the garden is the possibility in which there is a raven in the garden instead, and suppose that Ann's evidence eliminates the raven possibility. Then (1) is true in c_{easy} . (One might think of c_{easy} as a context in which the question under discussion is whether there is a goldfinch or a raven in the garden, and one might think of Ann as someone in position to answer this easy question.) But suppose that in context c_{hard} there is a relevant possibility in which there is a canary in the garden instead, which Ann's evidence does not eliminate. Then (1) comes out false in c_{hard} . (One might think of c_{hard} as a context in which the question under discussion is whether there is a goldfinch or a canary in the garden, and one might think of Ann as someone not in position to answer this hard question.)

It is worth mentioning that *Lewisian relevant alternatives theory* and *Lewisian contextualism* are independent. (This is why I have presented *Lewisian knowledge ascription* in conjunctive form, and why I will split the questions arising into two groups, in sections 30.2–30.3.) One can accept *Lewisian relevant alternatives theory* but reject *Lewisian contextualism*, simply by endorsing different rules of relevance. As mentioned above, one could even endorse rules that do not include any factors that can vary across contributors, thus rendering the resulting package invariantist. And likewise one can accept *Lewisian contextualism* but reject *Lewisian relevant alternatives*, simply by having a different account of the truth conditions for knowledge ascriptions. One could for instance add in a belief and/or justification condition. One could even endorse a semantics for knowledge ascriptions in which the relevance of alternatives played no role whatsoever, and relocate Lewis's rules of relevance to the realm of conversational pragmatics.¹⁰

Lewis then claims several virtues for the package that is *Lewisian knowledge ascription*, including two already mentioned: improving the connection between knowledge, belief, and justification (by treating the latter two as partial determinants of relevance rather than requirements for the truth of a knowledge ascription), and fitting our shifty intuitions for knowledge ascriptions (as seen in the treatment of (1) at c_{easy} and c_{hard}). With respect to fitting our shifty intuitions, subsequent empirical work looks to vindicate Lewis's intuitive idea that our intuitions about knowledge ascriptions are in fact sensitive to the contextually relevant alternatives.¹¹ (Though of course the invariantist can try to account for these intuitions in terms of conversational pragmatics or performance errors, instead of via the semantics.¹²)

But the main virtue which Lewis claims for *Lewisian knowledge ascription* is that of reconciling infallibilism with anti-skepticism. Thus Lewis (1996, 550) opens with a dilemma: “[W]e know a lot; knowledge must be infallible; yet we have fallible knowledge or none (or next to none). We are caught between the rock of fallibilism and the whirlpool of skepticism. Both are mad!” His account is primarily motivated as resolving this dilemma. The resolution is to grant the infallibilist her principle that knowledge requires the elimination of “every” relevant alternative, but understand this principle in a context-dependent way so that skeptical doubts (such as being a brain-in-a-vat) are not ordinarily relevant. Ordinarily such doubts do not come up for consideration at all, and so – in such ordinary contexts – knowledge ascriptions can still come out true. But in extraordinary contexts when such doubts come unto consideration, knowledge ascriptions go false, as he (1996, 551) explains:

Maybe epistemology is the culprit. Maybe this extraordinary pastime robs us of our knowledge. Maybe we do know a lot in daily life; but maybe when we look hard at our knowledge, it goes away. But only when we look at it harder than the sane ever do in daily life; only when we let our paranoid fantasies rip. That is when we are forced to admit that there always are uneliminated possibilities of error, so that we have fallible knowledge or none.

(Though the reader should keep in mind that when Lewis speaks of epistemology as robbing us “of our knowledge,” he really means something metalinguistic. He means that engagement in epistemology can create a type of context in which ‘know’ expresses an extremely demanding relation.)

Putting this together: Lewis offers an account of knowledge ascription – *Lewisian knowledge ascription* – which conjoins a distinctively metalinguistic and elegant form of relevant alternatives theory – *Lewisian relevant alternatives theory* – with a distinctively detailed and powerful form of contextualism – *Lewisian contextualism*. Not for nothing has “Elusive Knowledge” earned a central place in the contemporary contextualist canon alongside other pioneering work such as that of Cohen (1988) and DeRose (1995), and inspired a range of subsequent approaches including those of Schaffer (2005), Blome-Tillmann (2009), and Ichikawa (2011).

30.2 Epistemic Questions

So far I have spelled out Lewis’s theory via *Lewisian knowledge ascription*, which conjoins *Lewisian relevant alternatives theory* and *Lewisian contextualism*. But the theory is in some respects incomplete, and in other respects problematic. Without pretense of covering every possible concern, let me just pose some questions that any friend of contextualist relevant alternatives theory needs to address. I will begin with three questions which might be thought of as *epistemic questions*, insofar as they primarily concern *Lewisian relevant alternatives theory*. (The *semantic questions*, to be discussed in section 30.3, will be those that primarily concern *Lewisian contextualism*.)

30.2.1 Hyperintensionality

One obvious question is whether a framework of possibilities (the sort of things that might or might not be relevant, on Lewis’s approach) provides a sufficiently *fine-grained* framework for a hyperintensional topic like epistemology. Assuming that $5 + 7 = 12$ and Fermat’s Last Theorem are both necessary truths holding at every possibility, *Lewisian Knowledge* entails that ‘s knows that $5 + 7 = 12$ ’ and ‘s knows that Fermat’s Last Theorem holds’ are both true, for any context c . This is counterintuitive. It seems that in most contexts we would affirm, of a normal human adult Ann who is not a mathematician:

- (2) Ann knows that $5 + 7 = 12$.

But we would reject:

- (3) Ann knows that Fermat’s Last Theorem holds.

So one might worry that *Lewisian relevant alternatives theory* is built on too coarse a framework.

Lewis (1996, 551–2) touches on the matter of hyperintensionality, accepting (in accord with his general views) that “there is only one necessary proposition” but allowing that this “known

proposition may go unrecognized when presented in impenetrable linguistic disguise, say as the proposition that every number is the sum of two primes.” But he immediately adds (before moving on to other matters): “These problems of disguise shall not concern us here. Our topic is modal, not hyperintensional, epistemology.”

Lewis’s parting comment is odd, since epistemology is evidently a hyperintensional topic. There is no such topic as “modal epistemology.” (Imagine someone whose account of knowledge failed to entail factivity, and who dismissed the worry on grounds that her topic was “nonfactive epistemology.”) And Lewis’s talk of disguise is odd as well, insofar as disguises are in no way integrated into his positive theory. (As it might have been if Lewis had taken knowledge to be a ternary relation between a subject, a proposition, and a guise.) So *Lewisian Knowledge* still entails that (2) is true iff (3) is true, for any context *c*. The seeming counterexample stands untouched.

Moreover, Lewis’s talk of disguise is unhelpful in *third-person cases* in which the ascriber can “see through the disguise.” For instance, if Beth is a math professor who has just carefully reviewed Wiles’s proof of Fermat’s Last Theorem, and Ann is a first-year student, it seems perfectly plausible that Beth might both affirm (2) and deny (3). Whatever is happening in this case, Beth (the ascriber) is clearly *not* failing to recognize the necessary proposition under an “impenetrable linguistic guise.” After all, she has just carefully reviewed Wiles’s proof.

Let me suggest an alternative answer on Lewis’s behalf, which begins in a more concessionary mode: epistemology is a hyperintensional topic. But this alternative makes no further concessions, and instead notes that hyperintensionality is a problem for everyone (*tu quoque*). No one has yet developed a widely accepted framework for handling hyperintensional matters. Perhaps one ultimately needs “structured propositions,” or “impossible worlds,” or guise relativity, or something else entirely. But the best way that one can proceed at present with hyperintensional matters is to offer an account in the usual intensional idiom and trust that – should a framework for hyperintensionality emerge – the account will prove smoothly extendible. *Lewisian Knowledge* proceeds in exactly this best way.

It should be stressed that the intensional framework – as developed by Lewis and others – represents our best developed semantic framework to date. *It is no objection to an account of knowledge ascriptions that it is well integrated with our best developed semantic framework to date.* It would be an objection if it could be shown that the success of the account was somehow crucially reliant on known failings of our best framework, and so could not possibly extend to any potential successor frameworks. But nothing like that has been shown.

Indeed, suppose one chooses to extend the usual space of “possible worlds” by tacking on an outer sphere of “impossible worlds.” Then the proposition expressed by ‘ $5 + 7 = 12$ ’ will differ from the proposition expressed by ‘Fermat’s Last Theorem holds’, insofar as there will be mathematically impossible worlds in which the former holds but not the latter (and worlds in which the latter holds but not the former, and worlds in which neither holds). Should such a framework emerge as the best framework for hyperintensionality, then relevant alternatives theory would extend perfectly smoothly. I don’t mean to suggest that this is the best framework for hyperintensionality (the jury is still out), or that relevant alternatives theory requires this to be the best framework for hyperintensionality (compatibility with frameworks has to be assessed case-by-case). But I take this example to show that relevant alternatives theory can extend very smoothly to at least some potential successor frameworks. And so I would conclude that, with respect to hyperintensionality, relevant alternatives theory is actually doing *as well as is currently possible*.

30.2.2 Truth Conditions

Even given a framework of possibilities, one might wonder whether the truth conditions for a knowledge ascription can be *equated* with the condition of the subject’s evidence eliminating the relevant

alternatives, as per *Lewisian relevant alternatives theory*. There are a number of reasons that might be given for doubting the equation, but I will focus on the problem of *cheap knowledge*, arising in contexts in which the relevant alternatives are eliminated too easily.¹³ In some cases of cheap knowledge the only relevant alternatives are *silly*. For instance, if Ann has never even heard of either David Lewis or of Princeton University, it might thereby seem false for anyone to say:

- (4) Ann knows that David Lewis was a professor at Princeton University.

But suppose that Ann is a rabid baseball fan and has full knowledge of everyone who has ever played shortstop for the Chicago Cubs. Then in a context *c* in which David Lewis's being a shortstop for the Chicago Cubs is the only relevant alternative to him being a professor at Princeton University, *Lewisian relevant alternatives theory* has it that (4) is true in *c*.¹⁴ In other cases of cheap knowledge the relevant alternatives are *empty* (either because there are no $\sim p$ -worlds, or because none happen to be relevant). For instance, in a context *c* in which no alternatives to David Lewis being a professor at Princeton University are relevant, *Lewisian relevant alternatives theory* again has it that (4) is true in *c*, without requiring Ann to have any evidence whatsoever. Where the relevant alternatives are empty everyone trivially satisfies the truth conditions given by *Lewisian relevant alternatives theory*.

One tempting idea for solving the problem of cheap knowledge is to tinker with *Lewisian contextualism*, so as to ensure that there are always substantive relevant alternatives. But given that there are any propositions that are true at all worlds and that require substantive evidence to be known (e.g. that Fermat's Last Theorem holds), this move is hopeless. Given that the proposition in question is true at all worlds, no amount of tinkering with the rules of relevance can make an alternative relevant, simply because there are no alternatives out there that might be made relevant.¹⁵

Let me suggest an alternative move on Lewis's behalf, which involves tinkering with *Lewisian relevant alternative theory* to add further requirements (in ways that admittedly sacrifice some of its elegance, and reintroduce some of the epistemic requirements that Lewis sought to banish). For instance, one might reintroduce a belief requirement, as well as a basing requirement, as per:

Lewisian relevant alternatives theory, modified: A sentence of the form '*s* knows that *p*' is true in context *c* iff (i) *s*'s evidence eliminates every not-*p* possibility relevant in *c*, and (ii) *s* believes that *p* on this basis.

For instance, if Ann has never heard of David Lewis or Princeton University, then even though Ann's evidence might eliminate every possibility in which David Lewis was not a professor at Princeton University in a given context *c* (either because *c* accords no relevant alternatives to this claim or silly alternatives), (4) will still come out false in *c*, simply because Ann will lack the relevant belief, or in more complicated cases will lack the right basis for this belief.¹⁶

Lewisian relevant alternatives theory, modified still allows that a subject can have *cheap evidence*. For instance, Ann is still accorded all the evidence required by (4) in certain contexts. This might seem itself problematic, though I think that the notion of evidence in play is partially a term of art. What *Lewisian relevant alternatives theory, modified* does is to prevent cheap evidence from automatically transmuted into cheap knowledge, simply by requiring more than evidence for knowledge.

30.2.3 Knowledge Itself

Even given a framework of possibilities and the equation of the truth conditions for a knowledge ascription with the condition of the subject's evidence eliminating the relevant alternatives, there remains a question of *knowledge itself*. When '*s* knows that *p*' is true in *c*, what – metaphysically

speaking – must hold in the world? (This question is obscured by Lewis’s simplifying device of speaking as if he were discussing knowledge itself: section 30.1.1. One must first see that Lewis is only discussing knowledge ascription, in order to see that Lewis has left open the status of knowledge itself.)

There are at least three main sorts of answers that Lewis might give concerning the metaphysics of knowledge. One answer – that of the *epistemic pluralist* – begins by positing a plurality of binary knowledge-like relations K_1 – K_n . What varies with context is then which epistemic relation K_j a given occurrence of ‘know’ denotes. This is to treat ‘know’ as having a contextually variable denotation, and the world as sporting a plurality of knowledge-like epistemic relations. But a second answer – that of the *epistemic monist* – posits a single ternary knowledge relation $Kspx$. What varies with context is which value for the third relatum x is expressed by an ‘ s knows that p ’ sentence at a given context. This is to treat ‘know’ as having an invariant denotation, and the world as sporting a single knowledge relation, albeit one that is more complex than usually thought. And a third answer – that of the *epistemic nihilist* – posits no knowledge-like relations whatsoever. This theory posits only information about which alternatives various evidence would eliminate. What varies with context is the connection between an ‘ s knows that p ’ sentence and the eliminative information it conveys. So in terms of the metaphysics of the knowledge relation, is Lewis a pluralist, monist, or nihilist about knowledge itself?¹⁷

I don’t believe that anything in Lewis’s discussion turns very heavily on the choice between pluralism, monism, and nihilism about knowledge itself. But it does seem to me that epistemic monism is independently the most plausible choice. Semantically, it fits the plausible idea that ‘know’ has an invariant denotation (this will be discussed further in section 30.3.2). And metaphysically, it fits the natural idea that there is a real unity to knowledge. In this respect the question of knowledge itself is merely a respect in which Lewis’s own account is incomplete, not a respect in which it is problematic. Indeed I am suggesting that this gap in Lewis’s account can be filled in a plausible way.

30.3 Semantic Questions

I have just considered three questions which might be thought of as epistemic questions, insofar as they primarily concern *Lewisian relevant alternatives theory*. I will conclude by considering four further questions for Lewis’s account, which might be thought of as *semantic questions*, insofar as they primarily concern *Lewisian contextualism*. (I would reiterate that these questions are just some questions that any friend of contextualist relevant alternatives theory should address, and are not intended to exhaust every possible concern.)

30.3.1 ‘Know’ Itself

Just as Lewis’s account falls silent on knowledge itself, it likewise falls silent on the semantics of ‘know’ itself. *Lewisian knowledge ascription* offers truth conditions for a complex expression, of the form ‘ s knows that p ’. But presumably – given a compositional approach to semantics – the truth conditions for complex expressions are to be derived from the meanings of their component parts. Moreover, not all knowledge ascriptions have the form ‘ s knows that p ’. For instance, there are also knowledge-*wh* constructions such as:

- (5) Ann knows what kind of bird is in the garden.

Lewisian knowledge ascription falls silent on these sentences. So what does Lewis think is the denotation of ‘know’ itself, which operates in all knowledge ascriptions (including sentences like 5), and from which *Lewisian knowledge ascription* should derive?

Lewisian contextualism describes the determinants of relevance, but leaves open how this connects to the denotation of ‘know’ itself. That said, there is a very natural treatment of the denotation of ‘know’ available to Lewis, which preserves the epistemic monist’s idea of an invariant denotation (section 30.2.3), while extending his contextualism to other knowledge ascribing sentences, including (5). That is:

$[[\text{know}]]^{c,w,t} = \lambda s \lambda p. s\text{'s evidence eliminates every possibility in which not-}p \text{ that is relevant in } c.$

On such a treatment, ‘know’ (relative to a context of utterance, and a world and time of evaluation¹⁸) invariably denotes a ternary relation between a subject, a proposition, and the contextually given relevant alternatives to the proposition (the contrasts), obtaining whenever the subject’s evidence eliminates these contextually given relevant alternatives.

Or if one preferred to operate with *Lewisian relevant alternative theory, modified* (adding in belief and basing requirements), then one would reach:

$[[\text{know}]]^{c,w,t} = \lambda s \lambda p. s\text{'s evidence eliminates every possibility in which not-}p \text{ that is relevant in } c, \text{ and } s \text{ believes that } p \text{ on this basis.}$

Equivalently, one could treat the relevant possibilities in which not-*p* as a contrast proposition *q*:

$[[\text{know}]]^{c,w,t} = \lambda s \lambda p. s\text{'s evidence eliminates } q \text{ (where } q \text{ disjoins the possibilities in which not-}p \text{ that are relevant in } c), \text{ and } s \text{ believes that } p \text{ on this basis.}$

The result would be a contrastivist semantics for ‘know’ itself, in the spirit of Schaffer (2004b, 2005, 2007), and as developed in Schaffer & Szabó (forthcoming).¹⁹

So the matter of the semantics of ‘know’ itself is a further matter on which Lewis’s account is incomplete but not necessarily problematic. Indeed I am suggesting that this gap in Lewis’s account can be also filled in a plausible way.

30.3.2 *Is ‘Know’ an Indexical?*

Even given a contextualist denotation for ‘know’ itself (as suggested in the previous section), one might worry whether contextualism is generally semantically plausible. Does such a denotation look like anything else we find in semantics? Or is it an *ad hoc* invention? Lewis offers a range of guiding analogies for his contextualist proposal. He (1996, 563–4) speaks of ‘know’ as being like an indexical pronoun, like a gradable adjective, and – perhaps most centrally for his semantics – as being like a quantificational determiner with a contextually restricted domain argument. But these guiding analogies are inequivalent, and moreover each analogy is individually questionable.²⁰

These guiding analogies are inequivalent, insofar as the leading treatments of indexical pronouns, gradable adjective, and quantificational determiners involve quite different structures. Indexical pronouns have no extra structure, but instead might be thought of (following Kaplan 1989) as having a context invariant character which is a function from contexts to contents. Gradable adjectives (at least on the treatment of Kennedy and McNally 2005) have additional structure in the form of a

degree argument, and associate with comparative and degree morphology, while quantificational determiners (on a view which traces back to Heim 1982) have quite different additional structure in the form of a restrictor argument.²¹

Moreover these guiding analogies are each individually questionable. Semantically, ‘know’ does not seem to behave like indexical pronouns such as ‘I’, ‘here’, and ‘now’, which are all obviously context sensitive in a way that ‘know’ is not, and which exhibit smooth tracking across contexts (e.g. we have no trouble tracking the referent of ‘I’ as speakers take turns in a conversation) in a way that ‘know’ does not. And syntactically ‘know’ does not seem to project either a degree argument or a restrictor argument, which (respectively) explain the context sensitivity of the semantics for gradable adjectives and quantificational determiners.

The resulting concern about the semantic plausibility of contextualism constitutes, to my mind, one of the main concerns with *Lewisian contextualism* (and contextualism generally). The friend of *Lewisian contextualism* might seek to defend one of the guiding analogies (for instance, see Blome-Tillmann 2008 and Ichikawa 2011). In my opinion the prospects for such a defense are grim, though I should acknowledge that this calls for further discussion. Or she might abandon the search for a guiding analogy, on grounds that our understanding of the context sensitivity of language is too primitive, and/or that the context sensitivity of ‘know’ might be *sui generis* (DeRose 2009, 13). But what is that, other than an admission that contextualism still looks – at least to the best of our limited current knowledge – to be implausible?

So I would suggest a different move on Lewis’s behalf, which is to find a better guiding analogy. Along these lines, Schaffer and Szabó (forthcoming) suggest that ‘know’ is best understood as an *A-quantifier*, where A-quantifiers traditionally include adverbial quantifiers, modal auxiliaries, and – on the classic views of Hintikka 1969 and Heim 1992 – attitude verbs. Rather than having a syntactically projected restrictor argument as per quantificational determiners (D-quantifiers), A-quantifiers pick up their restrictors at the semantic level, via a “modal base” argument interpreted via general discourse features (the question under discussion: section 30.3.3). A-quantifiers display a characteristic range of features, including: sensitivity to the question under discussion, association with focus, ability to take explicit ‘if’-clause as restrictors, and ability to coordinate domains with other A-quantifiers. Schaffer and Szabó then argue that ‘know’ displays these characteristic marks of being an A-quantifier, though it would only be fair to say that “the jury is still out” on whether an A-quantification model – or any other model – will prove viable.

30.3.3 Deriving Relevance

Even given a contextualist denotation for ‘know’ itself, augmented with a plausible guiding analogy, there remains room to question the specific list of rules that comprise *Lewisian contextualism*. One might question the details, by adding, subtracting, and/or revising a handful of the rules on the list Lewis provides.²² But at a deeper level – which I will focus on – one might question the entire project. For it seems that what Lewis is giving us in the end, however exactly the details are settled, is a somewhat ad hoc laundry list of rules specifically tailored to the epistemic case (Ichikawa 2011, 386–7). One might think – given that relevance plays systematic rules throughout semantics – that a theory of relevant alternatives in epistemology should be *derived* from a single general notion of relevance at work elsewhere in semantics.

Lewis sometimes speaks as if his rules of relevance should be derived from the context sensitivity of ‘every’. In this vein, Lewis (1996, 553) says that his requirement that “every” not-*p* possibility be eliminated is the source of his contextualism: “[W]e must attend to the word ‘every.’ What does it mean to say that every possibility in which not-*P* is eliminated? An idiom of quantification, like ‘every,’ is normally restricted to some limited domain.” But Lewis in fact makes no attempt to show

that any of his specific rules derive from the general context sensitivity involved with ‘every’, nor could he since many of his rules are tailored to the epistemic case.

Indeed – as Stanley (2005, 66) has argued – among the general features of the context sensitivity of quantificational determiners like ‘every’ is a natural capacity to shift freely between occurrences. Thus one finds sentences such as:

- (6) Every sailor waved to every sailor.

There is a natural reading of (6) in which the first domain of sailors and the second domain of sailors are entirely disjoint, for instance if the speaker is discussing a scenario in which two ships passed at sea and every sailor on the first ship waved to every sailor on the second ship. But among the specific features of the context sensitivity of ‘know’ that Lewis posits is a lack of such free shifting. Specifically, once skeptical scenarios come into play, subsequent knowledge ascriptions in that discussion are supposed to go false (as befitting what Lewis (1996, 551) calls the “irresistible” force of skepticism), which – given *Lewisian knowledge ascription* – requires the skeptical scenarios to remain fixedly relevant for as long as the discussion persists. Indeed Lewis (1979, 355) speaks of this in terms of a general (and somewhat mysterious) asymmetry in accommodation: “[F]or some reason raising of standards goes more smoothly than lowering.” And: “Because of this asymmetry, a player of language games who is so inclined may get away with it if he tries to raise the standards of precision as high as possible...”. So it seems as if the context dependence that Lewis himself posits for ‘know’ works very differently from the context dependence one finds for quantificational determiners like ‘every’.

I would suggest – as a way to derive the epistemically relevant alternatives from a single general notion of relevance at work elsewhere in semantics, while explaining the lack of free shifting – reading the epistemically relevant alternatives off *the question under discussion*. That is, I would suggest:

- Lewisian contextualism, modified:** A possibility is relevant at a context if it is a possible answer to the question under discussion in that context.

The question under discussion is a fairly orthodox semantic posit, posited as an entry on Lewis’s (1979) “conversational scoreboard” to reflect what is being addressed at that point in the conversation. This posit serves to explain diverse phenomena including felicitous topicalization, distant ellipsis, Gricean relevance (“speak to the question”), relevant alternative sets for contrastive focus, and domain restriction for A-quantification (Carlson 1983; Ginzburg 1996; Roberts 2004; see Schaffer and Szabó forthcoming: section 30.3.3 for further applications). *Lewisian contextualism, modified* would thus derive the epistemically relevant alternatives in way that fits a general notion of relevance at work elsewhere in semantics. Indeed this thesis is especially fitting for a view which models ‘know’ as an A-quantifier, as per the Schaffer and Szabó view (§3.2).

Lewisian contextualism, modified would also explain the lack of free shifting, since the question under discussion is a relatively stable discourse-level matter (in contrast with domain restrictors for D-quantifiers like ‘every’, which presumably have independent syntactic realizations that can freely be coordinated or not). Indeed there is independent evidence that A-quantifiers like modals do not permit free shifting (Stanley 2005, 73), as seen in the unacceptability of sentences such as:

- (7) Ann can speak Finnish and Ann can only speak English.

If the question under discussion could freely shift from *whether Ann has the capacity to learn Finnish if she studies for years*, to *whether Ann in fact has learned any Finnish*, then (7) should be acceptable. The unacceptability of (7) thus provides evidence against free shift for the question under discussion.

From a deeper perspective, *Lewisian contextualism, modified* fits a conception of our concept of knowledge as used to indicate *who can answer the question*. Recall:

(1) Ann knows that there is a goldfinch in the garden.

Whether (1) is true in a given context seems to depend on whether the question under discussion is the easy question of *whether there is a goldfinch or a raven in the garden*, or the hard question of *whether there is a goldfinch or a canary in the garden*. This fits the intuitive shiftiness of knowledge ascriptions (section 30.1.3), Schaffer's (2005, §1) conception of the role of knowledge as *fingering answerers*, and Hookway's (1996, 7) insight into the general role that knowledge plays in *evaluating inquiry*: "The central focus of epistemic evaluation is...the activity of inquiry...When we conduct an inquiry... we attempt to formulate questions and to answer them correctly."

Lewisian contextualism, modified, however, does not support Lewis's *Belief* or *Actuality* rules, since these considerations are not in general required to be factored into possible answers to the question under discussion. As a result, *Lewisian contextualism, modified* cannot sustain Lewis's idea of shunting the justification, truth, and belief conditions into considerations of relevance (section 30.1.2), and so would require something like *Lewisian relevant alternatives theory, modified* instead (section 30.2.2). Indeed it would require a further modification that included a truth condition, such as:

Lewisian relevant alternatives theory, re-modified: A sentence of the form '*s* knows that *p*' is true in context *c* iff (i) *s*'s evidence eliminates every not-*p* possibility relevant in *c*, (ii) *s* believes that *p* on this basis, and (iii) *p* is true.

In this respect it seems to me that *Lewisian contextualism, modified* is generating useful constraints on the shape of the accompanying relevant alternatives theory, and can thereby help the friend of contextualist relevant alternatives theory settle the question of whether to treat a given factor via the contextualist aspect of her theory, or via the relevant alternatives aspect of her theory.

30.3.4 How Did He Do It?

I began my statement of Lewis's account (section 30.1.1) by noting that he states his account in terms that he later reveals to be an incorrect simplification, in speaking of when *s* knows that *p*, instead of the contexts in which '*s* knows that *p*' counts as true. Though Lewis presents this as something of a harmless simplification (which the careful reader can correct if she likes), I think that Lewis's simplified statement actually raises a serious and underappreciated problem for him, concerning how it could possibly have gone well (by Lewis's own lights).

The problem is that Lewis is presenting his theory in a philosophical context in which skeptical scenarios are explicitly under attention and therefore relevant (by his irresistible rule of *Attention*).²³ By *Lewisian knowledge ascription* (plus some plausible assumptions about the uneliminability of such skeptical scenarios, which Lewis himself embraces) this entails that virtually all knowledge ascriptions made in Lewis's paper are false. Yet Lewis's simplified object language exposition involves various positive claims of knowledge. So doesn't the theory itself predict that the simplified

presentation Lewis opts for should fail? That is, doesn't the theory itself predict that the reader – or at least the accommodating reader, who allows Lewis to make skeptical scenarios salient – should be plunged into a skeptical context and thereby deem most if not all knowledge claims to be false? And doesn't this conflict with the fact that Lewis makes positive knowledge claims in such a context, and that many if not all of these positive knowledge claims still seem true in the skeptical context Lewis has created?

Lewis (1996, 566–7) sees the problem: “Don't you smell a rat? Haven't I, by my own lights, been saying what cannot be said?” and “Does not my story deconstruct itself?” His (1996, 566) reply is that he has “bent the rules,” but was able to do so by relying on “the cardinal rule of pragmatics, which overrides every one of the rules I mentioned: interpret the message to make it make sense – to make it consistent, and sensible to say.” (This is the portion at the close of the paper at which Lewis mentions that he has offered a simplified and potentially misleading presentation; it is also the portion – alluded to in section 30.1.2 – at which Lewis renders all of his rules of relevance resistible.)

But Lewis's reply – and the general idea that one can bend the rules in this way via “the cardinal rule of pragmatics” seems to conflict with certain of Lewis's key tenets, especially Lewis's (1996, 561) motivating puzzle: “how can it be, when his conclusion is so silly, that the sceptic's argument is so irresistible?” For now Lewis's “cardinal rule” is making it possible to get away with ordinary knowledge ascriptions in a context in which skeptical doubts are relevant (indeed that is exactly what Lewis is getting away with). Thus this response seems to license what DeRose (1995, 28) labels “abominable conjunctions” such as “I don't know that I'm not a bodiless (and therefore handless) brain-in-a-vat, but I still know that I have hands.” Such abominable conjunctions should seem felicitous enough if one can just “bend the rules” when one reaches the second conjunct.²⁴

It seems to me that – by the lights of most plausible contextualisms, or at least those that are skeptic-friendly in allowing the skeptic to succeed in the contexts she creates – Lewis's presentation simply ought not to have succeeded. The theory predicts that the reader will judge Lewis's positive knowledge claims to be false, and so regard his discussion as an outright failure.

One option for the contextualist is to be less skeptic-friendly, and allow ordinary knowledge ascriptions to still count as true even when skeptical scenarios are relevant. This option certainly counts as a departure from Lewis's approach. And worse, on this option the contextualist may need to abandon her claim to reconcile ordinary knowledge with skeptical doubt, and may need to withdraw her claim to meet Lewis's (1996, 550) primary motivation of steering between “the rock of fallibilism and the whirlpool of skepticism.”²⁵

A second option for the contextualist is to chalk this up to *compartmentalization*. In this vein Lewis (1996, 565) imagines two epistemologists on a bushwalk, whose discussion of skepticism might be interspersed with claims that they know where they are going, or that they know which sort of viper slithers yonder, with little to no “relevance leakage” between these discussions. I think that compartmentalization is plausible for this sort of case, but would only add that Lewis's discussion in “Elusive Knowledge” is a single extended and unified discussion (uninterrupted by the appearance of lethal wildlife or other conversation-stoppers) for which compartmentalization is inappropriate. Indeed, if the contextualist allows for compartmentalization to occur easily within a single unified discussion, she will be in danger of re-licensing abominable conjunctions so long as one can just “switch compartments” between the conjuncts.

Though I am sympathetic with the broad outlines of Lewis's contextualist relevant alternatives theory, I find it hard to avoid thinking that the success of Lewis's own simplified presentation is actually a deep embarrassment for his contextualism, one for which an explanation remains elusive.

Acknowledgments

Thanks to Michael Blome-Tillmann, Stew Cohen, Keith DeRose, Jonathan Ichikawa, and Rae Langton.

Notes

- 1 For instance, both Langton (2004) and Schaffer (2004c) argue that Lewis's contextualism does not fit his skeptical worries about properties in "Ramseyan Humility," since his contextualism should dissolve these skeptical worries.
- 2 As Lewis (1973, 92) clarifies, in reply to Quinean concerns that counterfactuals are too shifty for serious treatment: "I am not one of those philosophers who seek to rest fixed distinctions upon a foundation quite incapable of supporting them. I rather seek to rest an unfixed distinction upon a swaying foundation, claiming that the two sway together rather than independently."
- 3 One might think that these topics are so closely connected as to never differ. For one might think that, by disquotation, 's knows that p' is true in c iff s knows that p. But disquotation so formulated fails for contextually sensitive terms (for instance, '[I am Schaffer' is true iff I am Schaffer] can be false if I am the speaker and you are the disquoter). And by Lewis's contextualist lights, 'know' is a contextually sensitive term.
- 4 In this respect, Lewis's account differs from most other classic relevant alternatives theoretic accounts (cf. Austin 1946; Goldman 1976; and Dretske 1981), which really are object language accounts of when s knows that p.
- 5 By way of comparison, Goldman (1976, 772) starts from the idea that "a person is said to know that p just in case he *distinguishes* or *discriminates* the truth of p from relevant alternatives." But he (1976, 785–6) ultimately embeds this condition in clauses (3.c.i) and (3.c.ii) of a vastly more complicated account, and one which is moreover restricted to the special case of non-inferential perceptual knowledge of an object having a property.
- 6 The list of expanders (/contractors) is the same as the list of irresistibles (/resistibles). This seems largely accidental. As far as I can tell, Lewis has no deep theoretical reasons for rejecting resistible expanders, and no deep reason for rejecting irresistible contractors, so long as these can be insulated from conflict with whatever irresistible expanders there might be.
- 7 For instance, Lewis leaves some of the details of *Resemblance* open. As Lewis notes, *Resemblance* is in danger of overgenerating relevant possibilities in a way that would lead to rampant skepticism. After all, if resemblance with respect to the subject's evidence is salient, then – at least on an "internalist" conception of evidence on which the subject has the same evidence when envatted as when embodied (a conception Lewis endorses on p. 556) – skeptical scenarios will be relevant whenever the subject's evidence is salient. So Lewis concludes that *Resemblance* needs some yet-to-be-specified restriction.
- 8 The full story: *Actuality* ensures that the subject s's world is relevant. Lewis's conception of elimination (1996, 553) entails that s's evidence can never eliminate her actuality. Together these ensure that if s is at a not-p world, then s's world will be relevant (by *Actuality*) and uneliminated (by the conception of elimination), which then suffices (by *Lewisian Knowledge*) for 's knows that p' to be false in all contexts.
- 9 Against a justification requirement, Lewis (1996, 551) asks: "What (non-circular) argument supports our reliance on perception, on memory, and on testimony?" He also talks about knowledge supported by forgotten reasons, and adds: "[W]e know the name that goes with the face, or the sex of the chicken, by relying on subtle visual cues, without knowing what those cues may be." Against a belief requirement, Lewis (1996, 556) invokes cases like that of Radford's (1966) timid student, who gives the right answer but feels no confidence in what she says.
- 10 Schaffer (2004a) plays with this idea on behalf of the skeptic, with the idea being that knowledge ascriptions, though virtually all *false*, still may be *felicitous* in ordinary contexts on the model of felicitous hyperbole. On this treatment Lewis's rules of relevance are re-understood as felicity conditions on hyperbole.

- 11 In this vein, Schaffer and Knobe 2012 (replicated by Buckwalter 2014) present a range of experimental results confirming a strong, stable, and unified pattern of intuitive sensitivity to the contextually relevant options.
- 12 For pragmatic accounts, see Rysiew 2001 and Brown 2005. For performance error accounts, see Nagel 2008 and Bach 2010.
- 13 See Lihoreau (2008) for a version of this objection to Lewis.
- 14 Though it may be that Lewis's rule of *Relevance* will preclude there being such a context, by forcing some further relevant alternatives, such as Lewis being a professor at Harvard University (I thank Jonathan Ichikawa for discussion on this point).
- 15 Though on certain ways of handling hyperintensionality, such as by having an indefinitely extensible sphere of impossible worlds, every proposition can be accorded alternatives. So there is room to try to handle hyperintensionality and resolve the problem of cheap knowledge by (i) having an indefinitely extensible sphere of impossible worlds which generates alternatives for any proposition, and (ii) tinkering with the rules of relevance to ensure that certain alternatives are always relevant, e.g. by having an irresistible expander that renders the nearest alternative(s) relevant: ("The nearest not-*p* possibility(s) are always relevant.")
- 16 Additional motivation for inserting a belief and basing requirement comes from cases of *missed clues* (Schaffer 2001), in which the subject's evidence indeed eliminates the relevant alternatives, but the subject herself has no appreciation of this fact.
- 17 One can find strands in Lewis's thought that might seem to favor each of these three perspectives. Epistemic pluralism fits Lewis's (1986, 59–60) general picture of the abundance of (non-fundamental) properties and relations, there being "one of them for any condition we could write down, . . .". Epistemic monism fits Lewis's (pers. comm.; cf. Schaffer 2004b, 97) embrace of my contrastivism, which I had proposed as a monistic rival: "The only thing we disagree about is whether we disagree." And epistemic nihilism fits Lewis's (1996, 563) concluding discussion of "What is it all for?" where he says: "Ascriptions of knowledge . . . are a very sloppy way of conveying very incomplete information about the elimination of possibilities."
- 18 I insert world and time indices as per Kaplan's orthodox semantic framework. In fact Lewis (1980) preferred to operate with not just world and time indices, but additional indices for location and standard-of-precision.
- 19 For related versions of contrastivism, see Morton and Karjalainen 2003 and Sinnott-Armstrong 2004.
- 20 The worry about semantic plausibility represents perhaps the leading objection to contextualism in the current literature. Thus Stanley (2005, 47) – one of the main proponents of this objection – maintains: "[T]he alleged context-sensitivity of knowledge ascriptions has no other parallel among the class of uncontroversial context-sensitive expressions." And Blome-Tillmann (2008, 29) – though himself defending the idea that 'know' is an indexical – acknowledges: "Epistemic Contextualism . . . has fallen into considerable disrepute recently. Many theorists have raised doubts as to whether 'know' is context-sensitive, typically basing their arguments on data suggesting that 'know' behaves semantically and syntactically in a way quite different from recognized indexicals . . .". See Schaffer and Szabó (forthcoming, esp. §2) for further discussion.
- 21 Lewis (1996, 554) thinks of gradable adjectives in quantificational terms, via: "Just as *P* is known iff there are no uneliminated possibilities of error, so likewise a surface is flat iff there are no bumps on it. We must add the proviso: Psst! – except for those bumps that we are properly ignoring." But on leading views (Kennedy and McNally 2005) gradable adjectives are not quantifiers but degree-relative notions. What makes a surface flat is not the number of bumps it has, but rather its degree of bumpiness.
- 22 Along these lines, Cohen 1998 and Blome-Tillmann 2009 each propose tweaks to Lewis's rules.
- 23 Indeed these scenarios come up for discussion on the very first page of "Elusive Knowledge," when Lewis (1996, 549) instructs us to let our "paranoid fantasies rip – CIA plots, hallucinogens in the tap water, conspiracies to deceive, old Nick himself . . .". Discussion of such scenarios recurs throughout Lewis's discussion.
- 24 Stanley (2005, 67) argues that Lewis's theory licenses abominable conjunctions, on grounds that (i) Lewis treats 'know' as an indexical (cf. section 30.3.2), and (ii) indexicals are capable of freely shifting within a discourse. Stanley takes this to show that Lewis should not treat 'know' as an indexical. What I'm pointing

out is an entirely separate reason why Lewis is committed to abominable conjunctions, which Lewis could not block simply by rejecting the treatment of ‘know’ as an indexical.

- 25 Though skeptic-unfriendly versions of contextualism might still be said to reconcile ordinary knowledge with skeptical doubt, insofar as they can still be said to explain the mechanism by which the skeptic’s claims can count as true, even if that mechanism turns out to be seldom activated or to be often thwarted by other conversational pressures. Thus DeRose (1995, 6) says: “[T]he important point is to identify the mechanism by which the skeptic at least threatens to raise the standards for knowledge. Whether the skeptic actually succeeds against a determined opponent in so raising the standards is of little importance.” See DeRose (2009, ch. 4) for a more detailed discussion of options for the contextualist, including DeRose’s preferred “Gap view” (2009, 144–8), on which the skeptic and her opponent thwart each other from reaching conversational equilibrium, so that neither speaks truly (or falsely either). I am grateful to Keith DeRose for discussion on this point.

References

- Austin, J.L. (1946). Other minds. *Proceedings of the Aristotelian Society*, 20, 149–87.
- Bach, K. (2010). Knowledge in and out of context. In J.K. Campbell, M. O’Rourke & H. Silverstein (Eds.), *Knowledge and skepticism*, pp. 105–36. Cambridge, MA: MIT Press.
- Blome-Tillmann, M. (2008). The indexicality of ‘knowledge’. *Philosophical Studies*, 138, 29–53.
- Blome-Tillmann, M. (2009). Knowledge and presuppositions. *Mind*, 118, 241–94.
- Brown, J. (2005). Comparing contextualism and invariantism on the correctness of the contextualist intuitions. *Grazer Philosophische Studien*, 69, 71–100.
- Buckwalter, W. (2014). The mystery of stakes and error in ascriber intuitions. In J. Beebe (Ed.), *Advances in experimental epistemology*, pp. 145–74. New York: Bloomsbury.
- Carlson, L. (1983). *Dialogue games: an approach to discourse analysis*. Dordrecht: Reidel.
- Cohen, S. (1988). How to be a fallibilist. *Philosophical Perspectives*, 2, 91–123.
- Cohen, S. (1998). Contextualist solutions to epistemological problems: scepticism, Gettier, and the lottery. *Australasian Journal of Philosophy*, 76, 289–306.
- DeRose, K. (1995). Solving the skeptical problem. *Philosophical Review*, 104, 1–52.
- DeRose, K. (2009). *The case for contextualism: knowledge, skepticism, and context*, Vol. 1. Oxford: Oxford University Press.
- Dretske, F. (1981). The pragmatic dimension of knowledge. *Philosophical Studies*, 40, 363–78.
- Ginzburg, J. (1996). The semantics of interrogatives. In S. Lappin (Ed.), *The handbook of contemporary semantic theory*, pp. 385–422. Oxford: Blackwell.
- Goldman, A. (1976). Discrimination and perceptual knowledge. *Journal of Philosophy*, 73, 771–91.
- Hawthorne, J. (2004). *Knowledge and lotteries*. Oxford: Oxford University Press.
- Heim, I. (1982). The semantics of definite and indefinite noun phrases. PhD dissertation, University of Massachusetts-Amherst.
- Heim, I. (1992). Presupposition projection and the semantics of attitude verbs. *Journal of Semantics*, 9, 183–221.
- Hintikka, J. (1969). Semantics for propositional attitudes. In J.W. Davis, D.J. Hockney & W.K. Wilson (Eds.), *Philosophical logic*, pp. 21–45. Dordrecht: Reidel.
- Hookway, C. (1996). Questions of context. *Proceedings of the Aristotelian Society*, 96, 1–16.
- Ichikawa, J. (2011). Quantifiers and epistemic contextualism. *Philosophical Studies*, 155, 383–98.
- Kaplan, D. (1989). Demonstratives, an essay on the semantics, logic, metaphysics, and epistemology of demonstratives and other indexicals. In J. Almog, J. Perry & H. Wettstein (Eds.), *Themes from Kaplan*, pp. 481–563. Oxford: Oxford University Press.
- Kennedy, C. & McNally, L. (2005). Scale structure and the semantic typology of gradable predicates. *Language*, 81, 345–81.
- Langton, R. (2004). Elusive knowledge of things in themselves. In F. Jackson & G. Priest (Eds.), *Lewisian themes: the philosophy of David K. Lewis*, pp. 130–7. Oxford: Oxford University Press.

- Lewis, D. (1969). *Convention: a philosophical study*. Cambridge, MA: Harvard University Press.
- Lewis, D. (1979). Scorekeeping in a language game. *Journal of Philosophical Logic*, 8, 339–59.
- Lewis, D. (1980). A subjectivist's guide to objective chance. In R. Jeffrey (Ed.), *Studies in inductive logic and probability*, Vol. 2, pp. 83–132. Berkeley: University of California Press.
- Lewis, D. (1982). Whether report. In T. Pauli (Ed.), *Philosophical essays dedicated to Lennart Åqvist on his fiftieth birthday*, pp. 194–206. Uppsala: Filosofiska Studier.
- Lewis, D. (1986). *On the plurality of worlds*. Oxford: Blackwell.
- Lewis, D. (1988a). What experience teaches. *Proceedings of the Russellian Society*, 13, 29–57.
- Lewis, D. (1988b). Desire as belief. *Mind*, 97, 323–32.
- Lewis, D. (1996). Elusive knowledge. *Australasian Journal of Philosophy*, 74, 549–67.
- Lewis, D. (1999). Why conditionalize? In D. Lewis, *Papers in metaphysics and epistemology*, 403–7. Cambridge: Cambridge University Press.
- Lewis, D. (2001). Sleeping Beauty: reply to Elga. *Analysis*, 61, 171–6.
- Lewis, D. (2009). Ramseyan humility. In D. Braddon-Mitchell & R. Nola (Eds.), *Conceptual analysis and philosophical naturalism*, pp. 203–22. Cambridge, MA: MIT Press.
- Lihoreau, F. (2008). Relevant alternatives contextualism and ordinary contingent knowledge. *Disputatio*, 24, 281–94.
- Morton, A. & Karjalainen, A. (2003). Contrastive knowledge. *Philosophical Explorations*, 6, 74–89.
- Nagel, J. (2008). Knowledge ascriptions and the psychological consequences of changing stakes. *Australasian Journal of Philosophy*, 86, 279–94.
- Radford, C. (1966). Knowledge – by examples. *Analysis*, 27, 1–11.
- Roberts, C. (2004). Context in dynamic interpretation. In L. Horn & G. Ward (Eds.), *The Handbook of Pragmatics*, pp. 197–220. Oxford: Blackwell.
- Rysiew, P. (2001). The context-sensitivity of knowledge attributions. *Noûs*, 35, 477–514.
- Schaffer, J. (2001). Knowledge, relevant alternatives, and missed clues. *Analysis*, 61, 202–8.
- Schaffer, J. (2004a). Skepticism, contextualism, and discrimination. *Philosophy and Phenomenological Research*, 69, 138–55.
- Schaffer, J. (2004b). From contextualism to contrastivism. *Philosophical Studies*, 119, 73–103.
- Schaffer, J. (2004c). Quiddistic knowledge. In F. Jackson & G. Priest (Eds.), *Lewisian themes: the philosophy of David K. Lewis*, pp. 210–30. Oxford: Oxford University Press.
- Schaffer, J. (2005). Contrastive knowledge. *Oxford Studies in Epistemology*, 1, 235–71.
- Schaffer, J. (2007). Knowing the answer. *Philosophy and Phenomenological Research*, 75, 383–403.
- Schaffer, J. & Knobe, J. (2012). Contrastive knowledge surveyed. *Noûs*, 46, 675–708.
- Schaffer, J. & Szabó, Z.G. (2014). Epistemic comparativism: a contextualist semantics for knowledge ascriptions. *Philosophical Studies*, 168, 491–543.
- Sinnott-Armstrong, W. (2004). *Pyrrhonian skepticism*. Oxford: Oxford University Press.
- Sosa, E. (1986). On knowledge and context. *Journal of Philosophy*, 83, 584–5.
- Stanley, J. (2005). *Knowledge and practical interests*. Oxford: Oxford University Press.
- Unger, P. (1975). *Ignorance: a case for scepticism*. Oxford: Oxford University Press.