Epistemic comparativism: a contextualist semantics for knowledge ascriptions

Jonathan Schaffer · Zoltán Gendler Szabó

Published online: 12 May 2013
© Springer Science+Business Media Dordrecht 2013

Abstract Knowledge ascriptions seem context sensitive. Yet it is widely thought that epistemic contextualism does not have a plausible semantic implementation. We aim to overcome this concern by articulating and defending an explicit contextualist semantics for ‘know,’ which integrates a fairly orthodox contextualist conception of knowledge as the elimination of the relevant alternatives, with a fairly orthodox “Amherst” semantics for A-quantification over a contextually variable domain of situations. Whatever problems epistemic contextualism might face, lack of an orthodox semantic implementation is not among them.

Keywords Epistemic contextualism · Knowledge · Semantics · Context

Questioning is the cutting edge of knowledge; assertion is the dead weight behind the edge that gives it its driving force.... Information may be the body of knowledge, but questioning is its soul.... [A]sking a question means envisaging alternatives, and only one at most of these alternatives can really exist. (Collingwood 1924, pp. 78–79)

Knowledge ascriptions seem *context sensitive*. For instance, if Ann was wondering who stole the diamonds and then found Claire’s fingerprints on the safe, then it would seem true for Ann to say:

1. I know that Claire stole the diamonds

J. Schaffer (✉)
Rutgers University, New Brunswick, NJ, USA
e-mail: jonathan.schaffer@rutgers.edu

Z. G. Szabó
Yale University, New Haven, CT, USA
e-mail: zoltan.szabo@yale.edu

Yet if Ann was instead wondering what Claire stole and then found Claire's fingerprints on the safe, then it would seem false for her to say 1. After all, the fingerprint evidence bears on the question of who stole the diamonds, but not on the question of what Claire stole.

Epistemic contextualism is the thesis that things are as they seem: knowledge ascriptions are context sensitive. But contextualism has faced sustained criticism on several fronts. Perhaps the main criticism is that contextualism is semantically implausible. Contextualists seldom articulate an explicit compositional semantics for 'know,' but usually just draw analogies with indexical pronouns, gradable adjectives, and/or quantificational determiners. These analogies are at best partial. And so many philosophers have grown openly skeptical of epistemic contextualism. As Stanley (2005, p. 47) writes: "[T]he alleged context-sensitivity of knowledge ascriptions has no other parallel among the class of uncontroversial context-sensitive expressions." Call this *the semantic problem*.

We aim to solve the semantic problem by articulating and defending a contextualist semantics for 'know.' Essentially, we will argue that 'know' may be treated as an A-quantifier (alongside adverbial quantifiers and modals). Our semantics integrates a fairly orthodox contextualist conception of 'know' as requiring the elimination of the relevant alternatives, with a fairly orthodox "Amherst" semantics for A-quantifiers as ranging over a contextually variable domain of situations. We call this a *comparativist* semantics, since it treats 'know' as comparing a contextually constrained domain proposition Q with an explicit scopal proposition P , requiring that every situation at which Q is true is a situation at which P is true.

Overview: In Sect. 1 we characterize and motivate contextualism, and introduce the semantic problem. In Sect. 2 we review the stock contextualist analogies, and explain why they are at best partial. Doing so will help us establish desiderata for a plausible semantics, and lead us to an analogy with A-quantifiers. In Sect. 3 we sketch an Amherst semantics for the A-quantifier 'always,' which we then use as a template for our comparativist semantics for 'know.' In Sect. 4 we draw further on the analogy with A-quantifiers to find new forms of evidence for comparativism, and to show how comparativism meets our earlier desiderata for a plausible semantics. We conclude in Sect. 5 by considering the extent to which comparativism might generalize to other attitude verbs.

1 Epistemic contextualism and the semantic problem

1.1 Characterizing epistemic contextualism

What is epistemic contextualism? Roughly speaking, it is the thesis that knowledge ascriptions are context sensitive. Somewhat more precisely, we use 'epistemic contextualism' to mean the following (c.f. Stanley 2005, p. 16; DeRose 2009, p. 2):

(*Epistemic contextualism*) A single knowledge ascribing sentence can bear different truth values relative to different contexts of utterance, where this

difference is traceable to the occurrence of ‘know,’ and concerns a distinctively epistemic factor.

The first clause of *Epistemic contextualism* characterizes the form of context sensitivity at issue: variation in truth value by a given sentence across contexts of utterance. The remaining clauses try to ensure that this variation is arising for the right reason: not due to some other element of the sentence (perhaps all sentences contain names, descriptions, or other context sensitive elements), and not due to non-epistemic features of ‘know’ (such as tense and mood features). This definition of epistemic contextualism could perhaps use refinement, but should suffice for our purposes.

Epistemic contextualism is consistent with a wide range of underlying epistemic and semantic views. Epistemically, *Epistemic contextualism* requires that knowledge can be sensitive to features encoded in the context. This is neutral as to whether knowledge is sensitive to a range of contrasts (Austin 1946; Lewis 1996; Schaffer 2005a, inter alia), to a requisite level of evidence (Cohen 1988, 1999), or to a given distance through which one must track truth (DeRose 1995), inter alia. *Epistemic contextualism* is also neutral as to whether context sensitivity plays any role of interest to traditional epistemologists, such as helping to resolve skeptical paradoxes.

Semantically, *Epistemic contextualism* requires different truth values at different contexts. This is neutral as to whether the truth value differences arise from the expression of different propositions, or from the evaluation of a single proposition at different index coordinates (as per the *nonindexical contextualism* characterized in MacFarlane 2009). Even given different propositions, *Epistemic contextualism* is neutral as to whether the proposition expressed at a context is a purely semantic affair, or whether there are also pragmatic determinants of what is said (as per Sperber and Wilson 1986; Bach 1994; Carston 2002; Soames 2008, inter alia). *Epistemic contextualism* is also neutral as to whether such a propositional difference arises because ‘know’ is an indexical term expressing different contents in different contexts, or a relational term involving a potentially covert contextually sensitive argument (as per Schaffer 2004).¹

Putting this together, *Epistemic contextualism* is best understood as a dual constraint on epistemic and semantic theorizing. Epistemically, *Epistemic contextualism* requires a conception of knowledge that can be sensitive to features encoded in the context, while being silent on which features these might be. Semantically, *Epistemic contextualism* requires knowledge ascriptions to come out with the requisite context sensitivity, while being silent on the semantic mechanisms by which this might come to be. The semantic problem is the worry that there is no plausible way to satisfy this dual constraint.

We aim to solve the semantic problem, by showing that there is a plausible way to satisfy the dual constraint of *Epistemic contextualism*. We will work with the

¹ In Schaffer (2004) and Blome-Tillmann (2008, 2009), contextualism is specifically identified with the view that ‘know’ is an indexical term. The treatment of ‘know’ as a relational term is treated as a nearby alternative to contextualism. We are here using *Epistemic contextualism* broadly, to label a genus of which indexical and relational treatments of ‘know’ are both species.

epistemic idea that knowledge is sensitive to a range of contrasts, together with the semantic idea that the truth value differences arise from the expression of different propositions, in a purely semantic way, through a context sensitive situation domain argument associated with ‘know.’ We work with these ideas mainly because we think they are fairly plausible. But we also work with these ideas because we take them to be fairly orthodox, thus allowing us to rebut the common idea that the epistemic contextualist must hold semantically heterodox views. So we aim not merely to solve the semantic problem, but moreover to solve the semantic problem *in orthodox terms*.²

1.2 Motivating epistemic contextualism

Why believe *Epistemic contextualism*? *Epistemic contextualism* may be motivated in various ways, but perhaps the most direct motivation comes from paired cases in which a single knowledge ascribing sentence seems to shift truth-value across contexts. *Epistemic contextualism* is then the thesis that things are as they seem. Thus consider:

(*Who*) Claire has stolen the diamonds. Ann and Ben are wondering who stole the diamonds, and Ann finds Claire’s fingerprints all over the safe. So Ann says to Ben:

1. I know that Claire stole the diamonds

We take 1 to be naturally read as true in the context set by *Who*. After all, the fingerprint evidence identifies Claire as the person who stole the diamonds. The evidence gathered thus answers the question under discussion. Ann has successfully resolved the inquiry. But now consider 1 in a different context:

(*What*) Claire has stolen the diamonds. Ann and Ben are wondering what Claire stole, and Ann finds Claire’s fingerprints all over the safe. So Ann says to Ben:

1. I know that Claire stole the diamonds

We take 1 to be naturally read as false in the context set by *What*, given that Ann has no further evidence. After all, the fingerprint evidence does not identify the diamonds as what Claire stole. The evidence gathered does not answer the question under discussion. Ann has not successfully resolved the inquiry. And so, given that *Who* and *What* differ only over the question under discussion in the context under consideration, it seems that the question under discussion plays a role in truth evaluation.³

² We make no claim to provide the only solution to the semantic problem. Let a thousand flowers bloom. Another approach well worth considering is the *inquisitive semantics* developed by Groenendijk (cf. Groenendijk and Roelofsen 2009). Indeed at the 2007 Central APA, Groenendijk (“The Dynamics of Inquiry”) described a dynamic implementation of Schaffer’s (2007) question-relative view of knowledge. See also Aloni and Égré (2008).

³ The reader familiar with the literature might be surprised that we use “question shifting” cases like *Who/What*, as opposed to the more usual bank cases of DeRose (1992, p. 913) and airport cases of Cohen

By way of empirical confirmation, Schaffer and Knobe (2011) conducted a range of studies and found a unified pattern of intuitive shifts in knowledge ascriptions, triggered by question shifts and other shifts in the contextually relevant alternatives.⁴ In further work, Schaffer presented one hundred participants with a vignette that started with:

(*New*) Peter has just smashed the jewelry store window, grabbed the rubies in the display case, and fled the scene. Mary the police detective is now on patrol. By chance, she walks past the jewelry store. She sees the broken glass and can tell that a theft must have just taken place, but she has not yet determined who stole what. So she begins her investigation. She first finds and identifies Peter's fingerprints on the display case, and then she locates the security camera and recognizes Peter filmed in the act of smashing the window. So she says (quite loudly) to herself: "I have no idea what was stolen, but it was clearly Peter who did the stealing." Mary then leaves the scene to file a report.

Fifty of the hundred participants then saw the following continuation:

(*NewWho*) David, who lives just across from the jewelry store, has seen everything from his kitchen window. He has witnessed the theft, watched Mary's investigation, and overheard her concluding words. David is wondering who stole the rubies, and says to himself: "Mary knows that Peter stole the rubies."

The other fifty participants saw the following continuation, differing only in the question David is wondering about in the final sentence:

(*NewWhat*) David, who lives just across from the jewelry store, has seen everything from his kitchen window. He has witnessed the theft, watched Mary's investigation, and overheard her concluding words. David is wondering what Peter stole, and says to himself: "Mary knows that Peter stole the rubies."

Participants—regardless of which continuation they received—were then asked the extent to which they agreed with what David says. The mean agreement rates on a

Footnote 3 continued

(1999, p. 58). For the record, we agree with the usual contextualist line on bank and airport cases, but think that our question shifting cases make better flagship cases for three reasons. First, the standard bank and airport cases fail to form minimal pairs. They differ in both contextual factors and in what is at stake for the subject, and hence invite the "interest relative invariantist" reply that what is at stake for the subject can make a difference to what she knows (Fantl and McGrath 2002; Hawthorne 2004; Stanley 2005). The question shifting cases form minimal pairs involving no difference in what is at stake for the subject. Second, the intuitions alleged to hold in bank cases have proven empirically elusive (Feltz and Zarpentine 2010; May et al. 2010; Buckwalter 2010), while the intuitions in question shifting cases have been empirically substantiated (Schaffer and Knobe 2011). Third, it is not clear which contextual parameter the bank and airport cases are supposed to engage with, while the question-shifting cases clearly engage with the question under discussion, which is an independently motivated contextual parameter (Sect. 4.1).

⁴ Schaffer and Knobe ran studies on question shift cases analogous to *Who* and *What*, and also looked at knowledge-*wh* constructions and explicit 'rather than'-clauses. They (2011, Sect. 4) conclude: "We found—in accord with the contrastivist prediction—a unified pattern of shifts in responses across these three different ways of manipulating contrasts."

Likert scale from 1 to 7 (with 7 being “completely agree,” 4 being “neutral,” and 1 being “completely disagree”) were 5.05 for *NewWho* and 2.13 for *NewWhat*:

The difference is statistically significant, and the agreement rates straddle the midline, crossing from the “agree” to “disagree” side. Thus it seems as if changing the question under discussion can flip people from agreement to disagreement with a knowledge ascription.

In another study, this time with a within-subject design, Schaffer presented two hundred participants with *New* and then asked each of them two questions invoking explicitly different contrasts. One hundred participants were asked the extent to which they agreed with each of the following two knowledge-*wh* sentences (fifty participants saw 2 before 3, the other fifty saw 2 after 3):

2. Mary knows who stole the rubies
3. Mary knows what Peter stole

The other hundred participants were asked the extent to which they agreed with each of the following two knowledge sentences with clefts (fifty saw 4 before 5, the other fifty saw 4 after 5):

4. Mary knows that it was Peter that stole the rubies
5. Mary knows that it was rubies that Peter stole

The mean agreement rates on a Likert scale from 1 to 7 (with 7 being “completely agree,” 4 being “neutral,” and 1 being “completely disagree”) were as follows:⁵

The differences are statistically significant, the agreement rates straddle the middle, and no order effects were detected for which question was seen first. This looks like a strong, stable, and unified effect of shifting the contextually relevant contrasts.

⁵ In “2 before 3” the means were 5.9 and 2.36. In “2 after 3” the means were 5.3 and 2.38. In “4 before 5” the means were 5.54 and 2.14. In “4 after 5” the means were 5.52 and 1.82.

That said, there are various possible confounds to consider. For instance, with the original *Who* and *What* cases above, perhaps people assume that Ann and Ben have different background evidence (perhaps in *Who* people assume that Ann and Ben have already figured out that it was diamonds that were stolen). The *New* vignette is designed to address this concern by explicitly clarifying exactly what evidence Mary has. Or perhaps people tend towards *de re* readings in some cases but not others (perhaps in all the who-contrast cases, “the diamonds/rubies” tends to be read *de re*: c.f. Schaffer and Knobe 2011, Sect. 6.3). The within-subject design should help alleviate this concern, insofar as participants who get to consider both sentences in sequence are likely to read them in parallel ways. Or perhaps the relevant agreement levels can be explained purely in terms of conversational pragmatics or performance errors (see Schaffer and Knobe 2011, Sect. 6.3 for a discussion of the most promising attempts). We take the question shifting cases to reveal a strong, stable, and unified pattern of intuitive sensitivity to the contextually relevant options, but cannot pretend to have ruled out any prospect of an alternative explanation. We only claim that, pending an alternative explanation, question shifting cases provide *prima facie* evidence for *Epistemic contextualism*.

1.3 The semantic problem

Why not believe *Epistemic contextualism*? *Prima facie* evidence notwithstanding, *Epistemic contextualism* has faced sustained criticism on several fronts. Perhaps the most prominent line of criticism is that it is semantically implausible. This is the semantic problem, which we aim to resolve.⁶

Given that *Epistemic contextualism* is best understood as a dual constraint on epistemic and semantic theorizing, the semantic problem is the worry that there is no plausible way to satisfy this dual constraint (Sect. 1.1). We consider this a serious worry. That said, there are at least two ways to remain unconcerned. First, one might hold that linguistically unconstrained pragmatic factors concerning general rationality play a role in truth evaluation. Call this *pragmatic mediation*. If so then a given knowledge ascribing sentence might bear different truth values in different contexts, but not in a way subject to specifically linguistic constraints. No distinctively *semantic* problem would remain.⁷

But pragmatic mediation is a highly controversial thesis. Moreover, given the reliance on the vagaries of general rationality, it is hard to use pragmatic mediation to make specific predictions. Accordingly, we see little prospect in using pragmatic

⁶ The semantic problem is pushed in Schiffer (1996), Hawthorne (2004), Cappelen and Lepore (2005), and Stanley (2005), inter alia. As Blome-Tillmann (2008, p. 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...”

⁷ In this vein, Ludlow (2008) offers a contextualism “on the cheap,” by arguing that pragmatic factors play a role in lexical meaning via a process of *modulation*. And Stainton (2010) offers a version of contextualism by arguing that pragmatic factors play a role in propositional structure via a process of *free enrichment*.

mediation to explain the specific pattern of question shifts we displayed in Sect. 1.2. Thus we will for present purposes assume—if only to make matters hard for ourselves—that the specific pattern of question shifts displayed in Sect. 1.2 is traceable to some feature of logical form. Or at least, we will argue that there is a purely semantic implementation of *Epistemic contextualism* which can explain question shifting cases without *requiring* pragmatic mediation (or anything else unorthodox).⁸

A second way to remain unconcerned is to note that the anti-contextualist must prove a negative existential, namely that there is no plausible way to satisfy the dual constraint of *Epistemic contextualism*. Given that *Epistemic contextualism* is plausible (Sect. 1.2), and that there are a wide range of possible ways to satisfy its dual constraint (Sect. 1.1), the contextualist might regard semantic matters as an unresolved detail, while remaining confident that some plausible implementation will be found. In this vein, DeRose (2009, p. 9)—in the only book length defense of *Epistemic contextualism* to date—explicitly resists articulating a semantics:

There is another way to move beyond generic contextualism that I won't be engaging in at all. That is to provide an account of the kind of underlying semantic linguistic mechanism responsible for the varying truth-conditions that govern knowledge-attributing and -denying sentences.

He (2009, p. 13) notes a potential argument as “very threatening” to his view:

Perhaps by classifying the various different ways that different context-sensitive terms behave, someone can come up with what they have good grounds to think is a list of the various different ways for terms to be context-sensitive, together with specifications or predictions about how context-sensitive terms of each of the types will ‘behave’ in various ways. And then it could turn out that ‘knows’ doesn't fit any of these models, and that therefore we have good grounds for thinking contextualism is false.

But he immediately adds: “At this point, however, I've seen no argument worth worrying about.” He notes that any attempt to list types of context sensitivity may be incomplete, and may leave out the very type of context sensitivity exhibited by knowledge ascriptions.

We think that DeRose may be overly skeptical of our understanding of context sensitivity in natural language. There are some fairly well established models of context sensitivity, and so one may form some defeasible empirical judgment as to

⁸ A related division (from Cappelen and Lepore 2005, pp. 2–9): one might subscribe to the *radical contextualist* thesis that there is contextual sensitivity in every expression, or to the *radical invariantist* thesis that there is contextual sensitivity in almost no expressions (beyond the core indexicals and demonstratives), or to the more orthodox *moderate contextualist* thesis that contextual sensitivity is associated with less than all but more than almost none of the expressions (beyond the core indexicals and demonstratives). Those who posit pragmatic mediation are generally if not always radical contextualists. In the main text we are assuming the more orthodox moderate contextualist view, as this is the only view which renders *Epistemic contextualism* neither trivially true (as on radical contextualist views) nor trivially false (as on radical invariantist views). Our question is whether the moderate contextualist should consider ‘know’ to be a contextually sensitive expression. See Szabó (2006) for relevant discussion of moderate contextualism.

whether *Epistemic contextualism* fits any of these models. If not, that looks bad. Thus we are sympathetic to the method of argument offered by Stanley (2005, p. 51):

[W]hat I am in position to provide is a good inductive case that knowledge ascriptions are not context-sensitive in a distinctively epistemological way, as the contextualist would have it. In the realm of the empirical, a good inductive case is all we can expect.

We only disagree with Stanley's anti-contextualist conclusion. In any case, we do not wish to quarrel over which side has the burden of proof. We simply propose to go beyond DeRose by articulating an explicit semantics, in a way that we think will answer Stanley's concern about semantic plausibility.

2 Semantic analogies

Contextualists often draw analogies between 'know' and indexical pronouns, gradable adjectives, or quantificational determiners. We think that these analogies are at best partial, and that the semantic problem has arisen primarily because of the ways in which these analogies fail. In this respect we agree with many of the critics of *Epistemic contextualism*, who find these analogies poor and conclude that no plausible theory has yet been offered. We just think that the contextualist can do better.

We will now explain why we think that the extant analogies are at best partial. Though this section will be primarily critical, our purposes are ultimately positive, in two respects. First, we will use the final analogy with quantificational determiners to develop our own analogy with A-quantifiers, which will serve as a springboard to move beyond the level of analogy towards an explicit compositional semantic proposal (Sect. 3). Secondly, by explaining why the extant analogies are implausible, we hope to identify desiderata for plausibility which we can ultimately use to assess our own proposal (Sect. 4.5).

2.1 The analogy with indexical pronouns

Is 'know' like an indexical pronoun? When Cohen (1988, p. 97) introduces contextualism, he says: "[T]he theory I wish to defend construes 'knowledge' as an indexical. As such, one speaker may attribute knowledge to a subject, while another speaker denies knowledge to that same subject, without contradiction." Similar claims are made by DeRose (1992 pp. 920–921), Lewis (1996, p. 564), Neta (2003, p. 398), and Blome-Tillmann (2008, p. 31), all of whom propose to treat the context sensitivity of knowledge ascriptions on the pattern of the context sensitivity of claims with indexical pronouns like 'I' and 'you,' or at least on the more general Kaplanian (1989) pattern of expressions with a fixed character but a variable content.⁹

⁹ As will emerge when we review the extant analogies, many of the leading contextualists are on record as endorsing multiple analogies. Perhaps the best way to read these contextualists is simply as maintaining that there are multiple partial analogies to be drawn each of which is in some respects useful.

Yet there are at least two respects in which ‘know’ is unlike indexical pronouns, which have fueled anti-contextualist arguments. First and perhaps foremost, real indexicals seem to be *smoothly tracked across contexts*. This point shows up in a range of stock objections to epistemic contextualism, all of which are premised on an indexical model. For instance, one aspect of smooth tracking is the way we naturally adjust indexicals in *indirect speech reports* concerning different contexts (Cappelen and Lepore 2005, pp. 88–89). For example, when Ann utters:

6. There will be a party tomorrow

what Ann says cannot be homophonically reported unless the report happens to occur on the same day as Ann’s utterance. If the report happens the next day it will—at least if given in English¹⁰—require indexical adjustment:

7. Ann said that there is a party today

Yet we do not see any such adjustments for indirect speech reports with ‘know.’ Thus if Ann utters:

8. Ben knows that the party is at Claire’s house

it seems that 8 can be homophonically reported:

9. Ann said that Ben knows that the party is at Claire’s house

even if the question under discussion and other contextual factors have shifted (as long as the referents of the remaining terms are held fixed).

A second aspect of smooth tracking of indexicals across contexts is *the absence of cross-contextual disagreement*. For example, when Ann utters:

10. I am hungry

Ben—on noting that he himself is not hungry—will not thereby suffer any tendency to think that what Ann said is false, and will not thereby feel any inclination to disagree. Yet it seems clear that people who make divergent knowledge claims in different contexts do tend to regard themselves as disagreeing. They might engage in debate, and one side might even retract (MacFarlane 2005). Indeed, it is part of the standard contextualist treatment of skeptical arguments that those in skeptical contexts suffer a tendency to think that (true) knowledge ascriptions in ordinary contexts are false, and might even (wrongly) retract such true claims.

A third and closely related aspect of smooth tracking of indexicals across contexts is *the absence of widespread and persistent error*. For example, when Ann

Footnote 9 continued

We would not disagree with such a claim. Ultimately we just think that contextualists need to move beyond the analogies and articulate a semantics.

¹⁰ Some indexicals in languages such as Amharic (Schlenker 2003) and also Slave and Zazaki (Anand and Nevins 2004) may lack this feature. But smooth adjustment in indirect speech reports is at least a good diagnostic for English indexicals.

says 10, virtually everyone will understand that ‘I’ denotes Ann. We do not see much confusion about the referent of ‘I,’ at least given sufficient contextual information. But it is part of the standard contextualist treatment of skeptical arguments that those in skeptical contexts suffer a widespread and persistent tendency to think that ‘know’ still denotes what it denotes in ordinary contexts, which is why they retain an intuition of knowing that they have hands. This gets called *semantic blindness* (Hawthorne 2004, pp. 107–111). In this vein, DeRose (1992, p. 920) gives the analogy of people in different rooms who disagree over ‘Frank is in this room’—leading Schiffer (1996, p. 329) to quip: “If that’s the solution, what the hell was the *problem*?”¹¹

Leaving smooth tracking across contexts behind, a second relevant disanalogy is that indexicals can *shift freely within a discourse* (Stanley 2005, p. 66). There may be extra-linguistic physical constraints on the extent to which the contents of ‘here’ and ‘now’ can be expected to shift within a typical discourse (especially if the participants are stationary and their interaction is brief), but there do not seem to be any linguistic constraints. For example, if Ann and Ben are riding the bus together and looking out the window as their bus passes from a new apartment complex to a burned-out tenement, Ann might say:

11. I would not mind living here, but of course I would absolutely hate to live here

To put the point another way, there are no overall discourse constraints that link Ann’s first ‘here’ to her second in 11 (although the physical facts tend to keep the referents close).

But it is part of the standard contextualist treatment of skeptical arguments that there are overall discourse constraints that make it easy to expand the relevant alternatives, but difficult to contract them once they are introduced into the discourse (Lewis 1979, p. 355). Indeed, if the content of ‘know’ could freely shift, then one should be able to felicitously utter the kind of sentences that DeRose (1995, pp. 27–29) labels *abominable conjunctions*, such as:

12. Ann does not know that she is not a bodiless and therefore handleless brain-in-a-vat, but she does know that she has hands

12 would be felicitous if the first occurrence of ‘know’ could denote a skeptical content, and the second occurrence of ‘know’ could freely shift to a non-skeptical content (Stanley 2005, p. 67). But 12 does indeed seem abominable, and should not count as felicitous in virtually any natural context. This suggests that the content of ‘know’ (or any other shiftable aspect the contextualist might entertain) must be

¹¹ Another one of the many related aspects of smooth tracking involves *propositional anaphora* (Stanley 2005, pp. 54–55), where one refers back to “what was said.” If Ann says ‘I am hungry,’ Ben may (noting that he is not hungry, but appreciating that Ann feels differently) say ‘I am not hungry, but what Ann said is still true.’ But if Ann claims to know that Claire stole the diamonds, and skeptical doubts arise in the interim, Ben cannot say ‘Ann does not know that Claire stole the diamonds, but what Ann said is still true.’

sensitive to overall discourse constraints that preclude free shifting within a discourse. Indexicals pronouns lack this feature.

Two desiderata for a plausible contextualist semantics have emerged so far. First, the semantics should not predict smooth tracking across contexts. Second, the semantics should predict overall discourse constraints that preclude free shifting within a discourse. For these reasons, the analogy with indexical pronouns is at best partial. Indeed we sympathize with the anti-contextualist critic who sees the analogy with indexical pronouns offered and questions the semantic plausibility of the view.

2.2 The analogy with gradable adjectives

Is ‘know’ like a gradable adjective? A second common contextualist analogy is between ‘know’ and gradable adjectives like ‘tall’ and ‘flat.’ For instance, Cohen (1986, p. 580) writes: “Compare ‘know’ with a term like ‘flat.’ Attributions of flatness can vary in truth value depending on what standards are applied... This is just what I want to say about ‘know’ and standards of evidentness for defeaters.” Likewise Hawthorne (2004, p. 53) begins his objections to contextualism by noting: “Contextualism about knowledge is often introduced and developed by analogy with various comparative adjectives. I shall do the same.” Indeed, gradability seems clearly to extend to verbs such as ‘believe,’ ‘remember,’ and ‘regret’:

13. Ann completely/mostly/somewhat regrets that she went to the party
14. Ann regrets that she offended Ben less/just as much/more than she regrets that she insulted Claire

Moreover, the analogy with gradable adjectives does not seem liable to the objection from smooth tracking across contexts that plagued the analogy with indexicals (Sect. 2.1). Gradable adjectives are not always smoothly tracked across contexts, and can occasion spurious debates. For instance, imagine that Professor Jones (at some US university) is frustrated by how athletes earn so much more than professors, and complains:

15. Professors are not rich

It seems that she may be homophonically reported in any context:

16. Professor Jones said that professors are not rich

The global hunger activist, speaking in a context where all US professors count as rich, may well claim to disagree with Professor Jones.¹² This is not to insist that gradable adjectives exhibit exactly the same cross-contextual behavior as ‘know.’ For present purposes we can remain neutral on that score. It is only to say that it is not nearly as obvious as it was with indexicals that there is a major disanalogy in cross-contextual behavior.

¹² In this vein, Cohen (1999, pp. 78–79) notes how easily Unger (1975) is able to convince people of flatness skepticism, by noting that virtually everything (Kansas, a pancake, a pool table surface, etc.) has some bumps.

Though the second objection—free shifting without discourse-level constraints—does seem to remain. In this vein, Stanley (2005, pp. 57–58) gives the examples of:

17. That butterfly is small, and that elephant is small
18. That butterfly is large, but that elephant isn't large

Within both 17 and 18 there is a natural free shift of the degree argument between one appropriate for butterflies in the first clause, and one appropriate for elephants in the second. 18 can be true, for instance, without the butterfly needing to be larger than the elephant: the butterfly may just be large *for a butterfly*, and the elephant not large *for an elephant*.¹³ For this reason the analogy with gradable adjectives does not support the kind of overall discourse constraints that should make it easy for the skeptic to expand the relevant alternatives, but hard for later uses of 'know' in the discourse to contract them. And the analogy does not support the kind of overall discourse constraints needed to explain the infelicity of abominable conjunctions like 12:

12. Ann does not know that she is not a bodiless and therefore handless brain-in-a-vat, but she does know that she has hands

But in any case—leaving aside how the analogy with gradable adjectives fares with respect to our first two desiderata—'know' just does not look like a gradable verb. As Dretske (1981, p. 363) remarks: "Knowing that something is so, unlike being wealthy or reasonable, is not a matter of degree." He continues: "In this respect factual knowledge is *absolute*. It is like being pregnant: an all or nothing affair." In support of Dretske's contention, Stanley (2005, pp. 35–46) argues that 'know' does not combine properly with comparative and degree morphology. Returning to the examples of 13 and 14 which suggested that 'regret' is gradable, we find no similarly acceptable counterparts with 'know':

19. ?Ann partly/mostly/fully knows that she went to the party
20. ?Ann knows that she offended Ben less/just as much as/more than she knows that she insulted Claire

That said, we freely admit that the data is complicated and that there are special cases in which 'know' does seem to naturally combine with comparative and degree morphology (Dutant 2007):

21. Ann knows well enough/very well/full well that speeding is dangerous
22. Ann knows better than many/most/the average American that speeding is dangerous

Indeed, 'know' combines naturally with comparative and degree morphology in knowledge-*wh* constructions:

¹³ One could reject the claim that 17 and 18 each involve a shift in the degree argument, by positing the content *large for its own kind* for 'large' (DeRose 2008, p. 153). But if such an interpretation for 'large' were freely available we would expect 'That butterfly is larger than that elephant' to have a true reading as well, since the butterfly's degree of largeness for its own kind may exceed the elephant's degree of largeness for its own kind.

- 23. Ann knows how to play chess extremely well/quite well/a bit better than Ben
- 24. Ann knows better than anyone/very well/pretty well when to act tough

We speculate that 21–24 involve some sort of idiomatic or *ad hoc* interpretation.¹⁴ Indeed it seems that there is some opportunity to use comparative and degree morphology, even where the underlying notion is evidently not gradable. For instance, people do say things like:

- 25. That is the truest thing you've said this week

despite the plausible claim (*pace* Weatherson 2005) that truth is an absolute notion. And despite Dretske's point that being pregnant is an "all or nothing affair," we routinely describe those in the later stages of pregnancy as "very pregnant." What seems telling is that even with a capacity for *ad hoc* interpretation we *still* cannot find a way to salvage 19 or 20.

With genuine gradables we find not just an association with comparative and degree morphology, but a semantic interpretation of this morphology with respect to a fixed underlying *scale*.¹⁵ For instance, with 'tall' the comparative and degree morphology gets interpreted with respect to the height scale, with 'expensive' the price scale, and with 'flat' the smoothness scale. In the realm of gradable verbs, with 'believe' the comparative and degree morphology gets interpreted with respect to a credence scale, with 'remember' a clarity scale, and with 'regret' a depth scale. Thus claims about what one regrets can typically be paraphrased as claims about the depth of one's regret, *all the way up and down the scale*:

- 26. I regret deeply/to a fair degree/a bit that you will not come to my party
- 27. My regret that you will not come to my party is profound/significant/minimal
- 28. The depth of my regret that you will not come is high/middling/low

But claims about what one knows are *not* amenable to similar paraphrases. Thus consider:

- 29. ?Ann's knowledge that she went to the party is profound/significant/minimal
- 30. ?The degree of Ann's knowledge that she went to the party is high/middling/low

¹⁴ In 23 and 24 perhaps there is literal grading, but of something other than the knowledge state. For instance, it is natural to read 'Ann knows when to act tough' as a *generic*, with the comparative morphology interacting with a covert generic operator, so that 24 gets read as saying that Ann knows when to act tough in a wider range of situations than anyone else. Alternatively the interaction might involve the embedded question. See Pavese (manuscript) for discussion of these options.

¹⁵ The leading semantics for gradable adjectives treats them as context sensitive relational (non-indexical) expressions, taking a *degree argument* (Kennedy and McNally 2005; Kennedy 2007). On this account, 'Jack is tall' says (roughly) that the degree of Jack's tallness exceeds *d*, where *d* is the contextual standard for tallness on the height scale.

Overall we regard the limited interaction of ‘know’ with comparative and degree morphology (as in 21–24) as puzzling, and we regard the idea that ‘know’ is semantically interpreted with respect to a fixed underlying scale as improbable. We mark a sharp contrast between ‘know’ and ‘regret’ in this regard.

A further desideratum for a plausible contextualist semantics has emerged. The semantics should not predict smooth interaction with comparative and degree morphology, and should not allow paraphrases up and down an associated scale. Thus the analogy with gradable adjectives is at best partial. We sympathize with the anti-contextualist critic who sees the analogy with gradable adjectives offered and questions the semantic plausibility of the view.

2.3 The analogy with quantificational determiners

Is ‘know’ like a quantificational determiner? A third common contextualist analogy is between ‘know’ and quantificational determiners like ‘all’ and ‘every.’ On the version of contextualism that is rooted in a *relevant alternatives* approach to knowledge (Sect. 3.2), knowing that *p* requires eliminating *every* possibility in which *p* is false. Lewis (1996, p. 553) holds that this ‘every’ introduces context sensitivity: “[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.” This suggests that semantically speaking ‘know’ works like a quantifier, involving a context sensitive domain argument.¹⁶ In this vein, Ichikawa (2011, p. 385) aims to “rehabilitate a Lewisian contextualist approach to knowledge,” involving “context-sensitive quantifier domains.”

Yet the analogy with quantificational determiners seems equally liable to the objection from smooth tracking across contexts that plagued the analogy with indexicals. For instance, imagine that Ann, having brought some bottles of beer home from the store, and having just finished loading the final bottle into the fridge, declares:

31. All the beer is in the fridge

If Ben then goes to the store to buy more beer, and is wondering where the beer is in the store, he cannot homophonically report Ann via:

32. Ann said that all the beer is in the fridge

Moreover, there will be no felt disagreement between Ann and the store clerk, if the clerk tells Ben:

¹⁶ While Lewis draws an analogy with quantifiers, he (1996, p. 564) also draws an analogy with indexicals, and (1996, p. 554) with gradable adjectives: “Unger suggest an instructive parallel. 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.” Lewis maintains the analogy with gradable adjectives because he treats them as quantifiers. For instance, he treats ‘flat’ as a quantifier over bumps. But ‘flat’ isn’t really a quantifier. What makes a surface flat is its degree of bumpiness, not the number of bumps it has. At any rate, we treat the analogies with gradable adjectives and with quantifiers separately.

33. All the beer is on the shelf in aisle seven

And there will be no tendency to widespread and persistent error about what Ann has claimed in 31. Everyone will understand that Ann's talk of 'all the beer' only concerned the beer she had just brought home from the store. There is not so much confusion about the domains of quantificational determiners, at least given sufficient contextual information: people do not seem to suffer "semantic blindness" on this matter.

Moreover, the analogy with quantificational determiners seems equally liable to the concern about free shifting without discourse-level constraints, which plagued both the analogy with indexical pronouns and the analogy with gradable adjectives. In this vein, Stanley (2005, p. 60) provides the example:

34. Every sailor waived to every sailor

which can be true when every sailor on the ship waived to every sailor on the shore. No account seems available of the kind of overall discourse constraints that should make it easy for the skeptic to expand the relevant alternatives, but hard for later uses of 'know' in the discourse to contract them. And thus no natural account of the infelicity of abominable conjunctions like 12 seems available:

12. Ann does not know that she is not a bodiless and therefore handleless brain-in-a-vat, but she does know that she has hands

since if there were free shifting within a discourse it should allow 12 to be felicitous, with a free shift from the first skeptical conjunct to the second non-skeptical conjunct.

That said, the analogy with quantificational determiners at least does not seem liable to concerns about supporting comparative and degree morphology. This is not because 'all' and 'every' are not associated with a scale: the quantificational determiners do naturally form a scale, from 'all' down to 'most' down to 'few' down to 'none.' Indeed 'many' and 'few' can combine with intensifiers like 'very' and have comparative and superlative forms. Rather it is because 'all' and 'every' are lexically constrained to select the top of the salient scale, thus precluding paraphrases in terms of the degree of "allness" up and down a scale.

But in any case—leaving aside how the analogy with quantificational determiners fares with our first three desiderata—"know" just seems to lack the structure of a quantificational determiner. Quantificational determiners typically combine with *overt nominal expressions* to form quantifier phrases (e.g. 'all the beer'), where the nominal expression ('the beer') plays a crucial role in fixing the domain: the quantifier does not range over anything not in the extension of the nominal expression.¹⁷ Knowledge ascriptions lack this structure. They do not typically

¹⁷ We call the combination with nominal expressions "typical" because of cases of ellipsis ('Some expressed concerns'), quantifier float ('The students have all passed') and predicative uses ('The protestors were many'). Moreover, there are bare occurrences of quantificational determiners in which the role of the nominal is more subtle. For instance, a natural reading of 'This election could have two

combine with any expression that plays the semantic role of nominal expressions in constraining the domain of quantification.

Thus another desideratum for a plausible contextualist semantics emerges: the semantics should not require a typically overt expression combining with ‘know’ to constrain its context sensitivity. Typically the context sensitivity should come from either ‘know’ itself or some covert element. Thus the analogy with quantificational determiners is at best partial. We sympathize with the anti-contextualist critic who sees yet another misleading analogy and concludes that the semantic problem is serious indeed.

2.4 Towards a new analogy with quantificational adverbs

Is there a better analogy? We think that there is a more promising analogy in the neighborhood of quantificational determiners, involving adverbial quantifiers and modals.¹⁸ Following Bach, Kratzer, and Partee (Partee 1995), we take it that natural language quantification comes in at least two main forms. There is D-quantification, modeled on *determiners*, in which syntax more or less straightforwardly delivers a tripartite [*Quantifier*] [*Restrictor*] [*Scope*] structure. But there is also A-quantification, modeled on *adverbial* quantifiers (and *auxiliaries* and *argument-structure adjusters*), without any obligatory overt restrictor, and for which syntax alone may leave the restrictor argument in quantificational structure underdetermined. With A-quantification the tripartite structure may depend on additional non-syntactic factors, including discourse-level phenomena (Sect. 4.1).

The analogy with A-quantifiers is more promising than the analogy with D-quantifiers, in at least three main respects. First, A-quantifiers select their domains differently from D-quantifiers, without requiring any overt restrictor. Thus consider Quine’s (1966, pp. 90–92) famous example:

35. Tai always eats with chopsticks

Only relative to a very extraordinary context might 35 mean that every single situation in the universe is one in which Tai is eating with chopsticks. Rather in some natural contexts 35 means that Tai eats with chopsticks *when he eats with anything*, and in at least some other natural contexts 35 means that Tai eats with chopsticks *when he does anything with chopsticks*. In the former sort of context, ‘eats’ plays a domain constraining role and ‘with chopsticks’ does not, while in the second sort of context the situation is reversed. Thus A-quantifiers do not require any overt restrictor (just like ‘know’), while still involving contextually variable domain selection.

Footnote 17 continued

winners’ (the “two real contenders” reading) says neither that it is possible that this election has two winners, nor that two winners are such that this election could have them (Szabó 2011).

¹⁸ In the literature one occasionally sees an analogy with modals (cf. Schaffer 2005b, p. 126), but to our knowledge no one has drawn an analogy with quantificational adverbs. Stanley (2005, p. 32) includes quantificational adverbs on his initial list of context sensitive expressions, but never returns to this option when arguing that ‘know’ is unlike any entries on his list. Presumably this was because no one had seriously advocated an analogy.

The second respect in which the analogy with A-quantifiers is promising is that there is independent reason to believe that certain *attitude verbs* have an A-quantificational aspect. Indeed there is a classic treatment of attitude verbs as verbal quantifiers which traces back to Hintikka (1962, 1969), and is further developed by Heim (1992), Percus (2000), and von Stechow (2004), inter alia. Since no overt restrictor is required, on this classic treatment *attitude verbs* are A-quantifiers.¹⁹

The third—and perhaps most intriguing—respect in which the analogy with A-quantifiers is promising is that A-quantifiers seem to show a pattern of context sensitivity that matches the pattern seen in our question shifting cases for ‘know’ (Sect. 1.2). Thus consider a question shifting case for ‘always’:

(*WhoAlways*) Claire has stolen the diamonds. Ann and Ben are wondering who stole the diamonds, and Ann learns that there have been numerous recent diamond and ruby thefts and that Claire has been the thief every time. So Ann says:

36. Claire always steals the diamonds

We take 36 to be true in the context set by *WhoAlways*.²⁰ After all, Claire is the person who has stolen the diamonds in every case in which someone has stolen the diamonds. But now consider 36 in a different context:

(*WhatAlways*) Claire has stolen the diamonds. Ann and Ben are wondering what Claire stole, and Ann learns that there have been numerous recent diamond and ruby thefts and that Claire has been the thief every time. So Ann says:

36. Claire always steals the diamonds

We take 36 to be false in the context set by *WhatAlways*.²¹ After all, it is not the case that diamonds have been stolen in every case in which Claire has stolen something. A suggestive parallel sensitivity to the question under discussion thus emerges for ‘know’ and ‘always.’ We flag this as one final desideratum for a plausible contextualist semantics: the semantics should predict question sensitivity.

We have drawn an analogy between ‘know’ and ‘always,’ but of course we would not say that the analogy is perfect: ‘know’ is a verb, and it bears a subject argument. Ultimately contextualists need to move beyond the level of analogies, and articulate an explicit compositional semantics for ‘know’ which may be judged in its own right. We only offer the analogy with ‘always’ as a starting point.

¹⁹ We do not fully endorse the claim that attitude verbs are verbal quantifiers, for two main reasons. First, we allow that ‘know’ is not a quantifier but merely has a quantificational aspect (Sect. 3.4). Second, we remain neutral as to whether all attitude verbs have a quantificational aspect (Sect. 5). That said, we still take the verbal quantifier treatment to provide some precedent for our semantics.

²⁰ 36 sounds most natural in the context of *WhoAlways* if the subject bears focal stress: “CLAIRE always steals the diamonds.”

²¹ The intuition that 36 is false in the context of *WhatAlways* is at least equally strong if the subject bears focal stress: “CLAIRE always steals the diamonds.” Thus even those who think that intonation impacts the semantic interpretation of 36 (such as Herburger 2000) can accept that there is a single logical form judged true in the one context false in the other.

What we have on the table at this point—through considering the ways in which the extant analogies fail, and offering a new analogy—are five desiderata for a plausible contextualist semantics:

- A plausible contextualist semantics should not predict smooth tracking across contexts
- A plausible contextualist semantics should predict overall discourse constraints that preclude free shifting within a discourse
- A plausible contextualist semantics should not predict smooth interaction with comparative and degree morphology, and should not allow paraphrases up and down an associated scale
- A plausible contextualist semantics should not require a typically overt expression to constrain the context sensitivity
- A plausible contextualist semantics should predict question sensitivity

And so we not only have a starting point for our own semantics via the analogy with ‘always’ (Sect. 3), we also have some ways to assess our semantics (Sect. 4.5).

3 Comparativist semantics

We will now use the analogy with A-quantifiers to develop a semantics for ‘know.’ We will begin with a fairly orthodox contextualist semantics for ‘always’ and a fairly orthodox relevant alternatives epistemology, and then use the former as a template for a semantics that implements the latter. We remind the reader that we are not trying to defend a contextualist semantics for ‘always,’ nor defend relevant alternatives theory. We are merely trying to reveal how these approaches can be *integrated* into an orthodox semantic implementation of *Epistemic contextualism*, thereby resolving the semantic problem.

We propose that ‘always’ and ‘know’ both involve quantification over a contextually variable domain of situations, expressing quantitative *comparisons* between a domain proposition and a scopal proposition. Indeed the analogy we posit runs deeper: the quantifiers involved with both ‘always’ and ‘know’ are universal; they are both restricted lexically, contextually, and sometimes explicitly to a domain of situations; and their scope is invariably given by an overt clause. But we also posit three main disanalogies: ‘know’ takes an additional subject argument for the knower; the lexical meaning of ‘always’ contributes a factual restriction on the domain while the lexical meaning of ‘know’ contributes an evidential restriction; and ‘know’ may involve additional components beyond that of the quantitative comparison. Due to the final point of disanalogy we will remain neutral as to whether ‘know’ *is* an A-quantifier or merely has an A-quantificational *aspect*.

3.1 Comparativism for ‘always’

What is a plausible semantics for ‘always’? One of the leading treatments of A-quantifiers like ‘always’ is the “Amherst” semantics which treats ‘always’ as a

universal quantifier, quantifying over a contextually variable domain of situations.²² Essentially, a sentence such as:

37. Claire always steals the diamonds

is treated as saying (roughly) that every situation in the domain is a situation in which Claire steals the diamonds. Likewise, ‘sometimes’ is treated as an existential quantifier, again quantifying over a contextually variable domain of situations.

Situations obviously play a crucial role so it may be useful to clarify some of their metaphysical features and semantic roles. Following Kratzer (1989, pp. 612–615) we take situations (or “cases” or “possibilities”) as having the following three metaphysical features: they need not be actual; they are non-repeatable world-, time-, and place-bound items; and they stand in parthood relations such that every situation is part of exactly one mereologically maximal situation: *a world*. The reader would not go too far awry, at least for present purposes, to think of situations as parts of Lewisian worlds (Lewis 1986). And we take situations to play the following three roles in semantics (Kratzer 2009): they map the semantic values of situation-neutral contents (which we’ll label ‘propositions’) to truth values, depending on whether or not the content holds at the situation or not²³; they are obligatorily represented in logical form (as opposed to being unarticulated constituents, or mere parameters of an index), and of course they are what A-quantifiers quantify over. Thus, supposing that the semantic value of a simple sentence like ‘Claire steals the diamonds’ (relative to a context) is:

$$\llbracket \text{Claire steals the diamonds} \rrbracket^c = \lambda s. \text{Claire steals the diamonds} (s)$$

what ‘always’ does in 37 is (roughly speaking) to take in this denotation and check that every situation in the domain has the relevant feature (e.g. being a situation in which Claire steals the diamonds).

A-quantification over situations (so understood) is then restricted in three distinguishable ways. First, there is a lexical restriction, which for adverbial quantifiers is a restriction to *actual* situations.²⁴ More precisely, if an ‘always’-sentence is true of a situation *s* then the quantification is restricted to situations which are worldmates of *s*; or as we will put it, to situations which are consistent with the world of *s*.

Secondly, there is a contextual restriction to *relevant* situations. 37, for instance, should not be taken to invariantly express the incredible falsehood that every single actual situation is one in which Claire steals the diamonds. Rather, in some natural

²² We dub this “Amherst” semantics due to its roots at UMass-Amherst, including the work of Berman (1987), Kratzer (1989), and the dissertation of von Stechow (1994), *inter alia*.

²³ In this respect situations are playing a semantic role typically allotted to worlds. But this is a generalization of, rather than a deviation from, orthodoxy: worlds are just the mereologically maximal situations.

²⁴ Adverbial quantifiers are often felt to have a slight modal flavor. Thus consider ‘Cats are never good swimmers’ and ‘Mice always go for cheese.’ These seem to license counterfactuals about what would happen to this cat were she to try to swim, and what that mouse would do were he near cheese. We ignore this complication.

contexts only situations in which Claire steals something are relevant, and what is said is that in all actual situations in which Claire steals something, she steals *the diamonds*. In other natural contexts only situations in which Claire does something to the diamonds are relevant, and what is said is that in all actual situations in which Claire does something to the diamonds, she *steals* them. And in yet other natural contexts only situations in which someone steals the diamonds are relevant, and what is said is that in all actual situations in which someone steals the diamonds, *Claire* steals them. (This is just furthering the point we raised in Sect. 2.4, via 35.)

Thirdly, there is an *optional* explicit restriction that may be given by explicit phrases such as ‘if’- and ‘when’-phrases. Thus consider:

38. Claire always steals the diamonds when she is bored

We take ‘when she is bored’ to restrict the quantifier to the situations it describes. We don’t want to take a stand on the meaning of ‘when,’ or the fine points of difference in meaning between ‘if,’ ‘when,’ and related terms. But eliding over these differences and ignoring internal structure, we take a clause like ‘when Claire is bored’ to denote a proposition (which is thereby eligible to restrict a quantifier over situations), along the following lines:

$\llbracket \text{‘when Claire is bored’} \rrbracket^c = \lambda s. \text{ Claire is bored } (s)$

So we assume that the (simplified) logical form of 38 is:

$[\text{always}] [\text{(when) she}_i \text{ is bored}] [\text{Claire}_i \text{ steals the diamonds}]$

What ‘always’ quantifies over in 38 are situations restricted (inter alia) to those in which Claire is bored.

Given that the domain of situations is restricted in these three distinguishable ways, we find it most natural to suppose that there are three distinct arguments in logical form (which we label the ‘*L-restrictor*,’ ‘*C-restrictor*,’ and ‘*E-restrictor*’ arguments respectively).²⁵ While we are not committed to any particular representation, we prefer the following ternary branching structure:

$[\text{A-quantifier}] [\text{L-restrictor}] [\text{C-restrictor}] [\text{E-restrictor}] [\text{Scope}]$

This structure is parallel to a structure naturally adopted for sentences like ‘Everyone who came had a good time’ in which it seems natural to assume that the lexical restrictor (perhaps represented by the bound morpheme ‘-one,’ serving to constrain the domain to persons) and the contextual restrictor are incorporated into ‘Everyone,’ and separated from the optional explicit restrictor ‘who came.’

²⁵ By distinguishing among three different types of domain restriction at the level of logical form we are departing from orthodoxy. For example, von Stechow (1994)—a theory we otherwise closely follow—leaves the actuality constraint on the domain of adverbial quantifiers implicit and represents the contextual parameter whose value is responsible for restricting the domain to relevant situations as a mere index on the quantifier. (That aside, our semantics could quite easily be adapted to von Stechow’s framework.)

The *L-restrictor* of adverbs of quantification, L^{adv} , is given by:

$$L^{adv} = \lambda s \lambda s'. s' \text{ is consistent with the world of } s^{26}$$

We will assume that the *C-restrictor* and *E-restrictor* positions are occupied by variables whose values are propositions. We will argue that the value of the implicit contextual restrictor variable is provided through the question under discussion (Sect. 4.1), and we will assume that the value of the explicit restrictor variable is the proposition expressed by the conjunction of any explicit restrictive clauses with the trivial proposition true in all situations (e.g. $\lambda s. s = s$). (This last trivial conjunct for the *E-restrictor* ensures that we have an argument but a non-restricting one in sentences without any explicit restrictive clauses.)

So putting this together, we have a quantifier over situations subject to a tripartite restriction, and scoping over a proposition. The simple approach (complications will come shortly) is to require that the scopal proposition holds in *quantifier-many* of the situations in the domain, via the schema:

Where Q is an A-quantifier with domain restricted by R and scope P , QRP is true iff Q -many R situations are P situations

For the universal quantifier ‘always,’ given our approach to domain restriction, this simple approach yields:

$$\llbracket \text{always} \rrbracket^c = \lambda E \lambda P \lambda s. \text{ every } s' \text{ such that } L^{adv}(s, s') \text{ and } C(s') \text{ and } E(s'), \text{ is also such that } P(s')$$

‘Always’ is thus treated as a universal situational quantifier with three arguments: an explicit restrictor E (whose value is the proposition expressed by the conjunction of any explicit restrictive clauses with the trivial proposition), scopal material P (whose value is the proposition expressed by the clause in the scope of ‘always’), and of course a situation argument s . The semantic value of ‘always’ needs no separate argument for the lexical restrictor because that is fixed for all occurrences of ‘always’ (via consistency with the world of s , as per L^{adv}), and no separate argument for the contextual restrictor C because that is fixed by context. Unpacking, ‘always’ denotes a function from two propositions and a situation to a truth value. The first proposition is the domain proposition, which is compared to the scopal proposition P via the requirement that every situation in which the domain proposition is true is a situation in which P is also true. We thus label this a *comparativist* semantics.

For the existential ‘sometimes’ this simple approach would replace ‘every’ with ‘some’:

$$\llbracket \text{sometimes} \rrbracket^c = \lambda E \lambda P \lambda s. \text{ some } s' \text{ such that } L^{adv}(s, s') \text{ and } C(s') \text{ and } E(s'), \text{ is also such that } P(s')$$

²⁶ This formulation assumes a type theory which guarantees that the values of the s' variable are situations. Alternatively, we could build this into the lexical restriction: $L^{adv} = \lambda x \lambda x'. x$ and x' are situations and x' is consistent with the world of x . We will continue to use typed variables (as is customary), but this is just a convenience.

The generalization to other adverbial quantifiers such as ‘usually,’ ‘mostly,’ and ‘rarely’ is straightforward. The semantics also generalizes elegantly to other A-quantifiers such as modals just by changing the lexical restriction on the domain. When ‘always ϕ ’ is true of a situation s the quantification is restricted to worldmates of s ; when ‘must ϕ ’ is true of a situation s the quantification is not lexically restricted at all.^{27,28} The lexical restrictor of modals, L^{mod} , is thus the trivial relation:

$$L^{mod} = \lambda s \lambda s'. s = s$$

Thus assuming that ‘must’ expresses universal quantification, the simple approach yields:

$$\llbracket \text{must} \rrbracket^c = \lambda E \lambda P \lambda s. \text{some } s' \text{ such that } L^{mod}(s, s') \text{ and } C(s') \text{ and } E(s'), \text{ is also such that } P(s')$$

And assuming that ‘can’ is the dual of ‘must,’ expressing existential quantification:

$$\llbracket \text{can} \rrbracket^c = \lambda E \lambda P \lambda s. \text{some } s' \text{ such that } L^{mod}(s, s') \text{ and } C(s') \text{ and } E(s'), \text{ is also such that } P(s')$$

The difference between ‘must’ and ‘always’ need not involve anything more than differences in their lexical restrictions. Thus emerges an elegantly unified picture of the general semantics of A-quantifiers.

But that said, the canonical versions of the Amherst semantics for A-quantifiers incorporate two complications. First, returning to ‘always,’ it seems too strong to require that the scopal material must hold in all situations in the domain. For even when we constrain the domain of quantification for 38, say, to all situations in which Claire does something to the diamonds and is bored, we will presumably still find some situations surrounding Claire’s thievings during which she is not then engaged in stealing the diamonds. For instance, we might find a situation of her reaching for the diamonds (just before stealing them), or a situation of her tucking the diamonds into her pack (just after stealing them).²⁹ The standard fix—due to Berman (1987)—is to require, not that quantifier-many situations in the domain are situations in

²⁷ On this treatment only contextual and explicit restrictions play a restrictive role in modals. Other treatments are possible within a general A-quantificational framework. For instance, assuming a fixed counterpart relation, one might treat modals as bearing a lexical restriction to worldmates of counterparts of s , as per:

$$L^{mod*} = \lambda s \lambda s'. s' \text{ is consistent with the world of a counterpart of } s.$$

²⁸ It is of course an empirical question whether there is a univocal modal ‘must’ (as we assume in the main text), or rather a polysemous term with distinct epistemic and deontic meanings. In the latter case there will be distinct lexical restrictions associated with the distinct meanings. Along these lines, we note that epistemic and deontic modals are lexically distinguished in languages like Javanese and St’át’imcets (indeed in St’át’imcets the distinction between universal and existential force is contextually rather than lexically determined: Matthewson et al. 2005).

²⁹ It is not wholly obvious that this is the *wrong* result. Perhaps in contexts in which Claire’s tucking the diamonds into her pack (after stealing them) really are in the domain it should be false to say ‘Claire always steals the diamonds,’ since sometimes Claire is doing things with the diamonds beyond stealing them. So one could try to use context sensitivity to resist this problem. We take no stand on the matter here.

which the scopal material obtains, but rather that quantifier-many situations in the domain are *parts of* situations in the domain in which the scopal material obtains. Thus for instance the situation of Claire touching the diamonds, and the situation of Claire tucking the diamonds into her pack, are both parts of situations in the domain in which Claire steals the diamonds.

The second complication—included mainly to handle donkey anaphora—is to qualify the quantification involved to quantification over *minimal situations*.³⁰ A minimal P -situation is a situation s such that $P(s)$ and s has no proper part s^* such that $P(s^*)$. Incorporating these two wrinkles yields the canonical version of the Amherst semantics for A-quantifiers:

Where Q is an A-quantifier with domain restricted by R and scope P , QRP is true iff Q -many minimal R situations are parts of a P situation

The application to ‘always’ and the generalization to other adverbial quantifiers and modals is a straightforward extension of the pattern seen with the simple approach given above.

Since the two complications just sketched won’t play any significant role in our discussion, we will adopt the expedient of working with the simple treatment. So let’s use the simple treatment to compute a semantic value for:

38. Claire always steals the diamonds when she is bored

This is being thought of as essentially composed of:

[always] [L^{adv}] [C] [Claire_i is bored] [she_i steals the diamonds]

The semantic value of the adverbial quantifier relative to a context c , on the simple treatment, is:

$\llbracket \text{always} \rrbracket^c = \lambda E \lambda P \lambda s. \text{ every } s' \text{ such that } L^{adv}(s, s') \text{ and } C(s') \text{ and } (E(s'), \text{ is also such that } P(s'))$

The ‘when’-clause denotes a proposition which provides the E argument:

$\llbracket \text{Claire}_i \text{ is bored} \rrbracket^c = \lambda s. \text{ Claire is bored } (s)$

‘she_i steals the diamonds’ denotes a proposition which provides the P argument:

$\llbracket \text{she}_i \text{ steals the diamonds} \rrbracket^c = \lambda s. \text{ Claire steals the diamonds } (s)$

Since $L^{adv}(s, s')$ says that s' is consistent with the world of s , what results from composing these semantic values—our semantic value for 38—is the proposition:

$\lambda s. \text{ every } s' \text{ such that } s' \text{ is consistent with the world of } s, C(s'), \text{ and } \text{Claire is bored } (s'), \text{ is also such that } \text{Claire steals the diamonds } (s')$

³⁰ See Heim (1990) and especially Elbourne 2005 for further discussion. The guiding idea is to treat donkey anaphors via proxy Russellian definite descriptions with uniqueness implications relativized to situations (cf. Heim 1990, and especially Elbourne 2005). Since we doubt the semantic nature of uniqueness implications for definite descriptions (cf. Szabó 2000, 2005a, b), we are not so moved by such arguments. In any case, the issue is orthogonal to our present concerns.

This proposition will be true—plausibly enough—if and only if every actual and relevant situation in which Claire is bored is one in which she steals the diamonds.

3.2 Relevant alternatives theory

We are proposing to integrate the comparativist Amherst semantics for ‘always’ (Sect. 3.1) with an epistemology on which knowledge requires the elimination of the contextually salient alternatives. Let’s be more precise about the epistemology. We begin from the following sketch:

(*Relevant alternatives theory*) To know is to have evidence that eliminates the relevant alternatives, and to form a true belief on this basis³¹

For present purposes we will leave the notions of evidence, elimination, truth, belief, and basing as intuitive notions. We would just emphasize that we are *not* committed to any *reductive analysis*. It may be that some of these notions can only be explicated in terms of knowledge. It is no part of our conception of an adequate semantic clause that it must yield any sort of conceptual reduction.³²

We actually need much less in one respect, and a bit more in another, than *Relevant alternatives theory* includes. *Relevant alternatives theory* can be understood as incorporating the following evidential assumption:

(*Evidence*) Knowledge requires having evidence that eliminates the relevant alternatives

Also incorporated in *Relevant alternatives theory* are some fairly anodyne truth, belief, and basing assumptions:

(*Truth*) Knowledge requires truth

(*Belief*) Knowledge requires belief

(*Basing*) Knowledge requires that the belief be formed on the basis of the evidence

We need much less than *Relevant alternatives theory* includes, in that we do not need *Truth*, *Belief*, or *Basing*. We adopt these largely for the sake of definiteness and orthodoxy. These assumptions can be modified, or waived entirely, consistent with a generally contextualist semantics of the sort we will develop. Also inessential is the

³¹ Such an approach is rooted in Wittgenstein’s (1969) and Austin’s (1946) approaches to knowledge, and versions of it have been defended by Dretske (1970, 1981), Unger (1975), Goldman (1976), Stine (1976), Lewis (1996), Neta (2002), Schaffer (2005a), Rysiew (2006), and Ichikawa (2011), inter alia.

³² In this respect our project is perfectly compatible with *conceptual primitivism* about knowledge, as endorsed most prominently by Williamson (2000). Williamson himself is no contextualist (Williamson 2005), partly due to concerns about linguistic implementation connected to semantic blindness, and partly for unrelated concerns about the role of knowledge in practical reasoning (an issue which is beyond the scope of this discussion). His primitive notion of knowledge thus presumably has less structure than our notion, involving just a single proposition rather than a quantitative comparison between two. But one can equally be a conceptual primitivist about knowledge as per our comparativist conception.

absence of any further conditions for knowledge.³³ Indeed we do not even need the full strength of *Evidence*, but only the aspect connecting evidence to “the relevant alternatives.” The idea that the needed relation between the evidence and the relevant alternatives is that of *elimination* is inessential. A contextualist semantics could treat the relation between the evidence and the relevant alternatives in other terms. For instance, one could require sufficient but fallible comparative support for what is known over the alternatives, or some purely modal pattern of covariation between belief and truth across the relevant situations, inter alia.

Though we also need a bit more than *Relevant alternatives theory* includes, for *Relevant alternatives theory* is not yet contextualist. We need to add that the relevant alternatives are contextually variable.³⁴ Putting this together, what we need for our implementation of *Contextualism* is:

(Key) Knowledge requires having evidence that involves a contextually variable domain of situations

With *Key* we connect together the contextual variability of A-quantifiers with a contextualist version of relevant alternatives theory, as both concerning a contextually variable domain of situations.

Of course *Key* is only essential to our particular implementation of *Epistemic contextualism*. One could attempt to detail a contextualist semantics that did not work with a contextually variable domain of situations, but worked with some other contextually variable item.³⁵ We have no objection to the attempt, but would just note that such a semantics would not look anything like our semantics, and presumably could not claim to treat ‘know’ as anything like an A-quantifier, given that the contextual variability in A-quantifiers concerns the domain of situations quantified over.

3.3 Comparativism for ‘knows’: the evidential aspect

We are now ready to integrate the comparativist Amherst semantics for ‘always’ with a contextualist version of relevant alternatives theory. We begin with *Evidence*,

³³ Of particular interest is the prospect of adding some *non-accidentality* condition. See Cohen (1998) and Heller (1999) for arguments that the proper solution to the Gettier problem does not come from the alternatives relevant to the attributor, but from considerations of the subject’s relation to the truth. One might also add a condition requiring that the subject has at least some evidence against the irrelevant alternatives (Dretske 1981, p. 373; Schaffer 2005a, p. 258).

³⁴ This is a substantive addition, since anti-contextualists of various stripes have in fact endorsed *Relevant alternatives theory*. Thus Unger (1975)—while defending skeptical invariantism—identifies the relevant options with the entirety of logical space. Dretske (1981, p. 377; cf. Dretske 1991, p. 192; Bach 2005, pp. 80–84)—defending a non-skeptical form of invariantism—identifies the relevant alternatives in a more restricted but still objective way, as those with non-zero objective chance. And Stanley (2005, p. 85; cf. Hawthorne 2004, pp. 158–161)—defending interest relative invariantism—allows for an implementation of his view in terms of relevant alternatives, where practical facts about the subject’s environment play a role in determining what is relevant.

³⁵ For instance, Cohen (1999, p. 59) works with contextual variability as to the strength of reasons sufficient to justify.

understood in a contextualist way as per *Key*. Our (simplified) comparativist semantics for ‘always’ will serve as a template:

$$\llbracket \text{always} \rrbracket^c = \lambda E \lambda P \lambda s. \text{ every } s' \text{ such that } L^{adv}(s, s') \text{ and } C(s') \text{ and } E(s'), \text{ is also such that } P(s')$$

Abstracting away from L^{adv} —the lexical restriction to actual facts characteristic of adverbial quantifiers—we reach the following (simplified) schema for universal A-quantifiers:

$$\llbracket A-Q \rrbracket^c = \lambda E \lambda P \lambda s. \text{ every } s' \text{ such that } L\text{— and } C(s') \text{ and } E(s'), \text{ is also such that } P(s')$$

Our proposed contextualist implementation of *Evidence* (*Truth*, *Belief*, and *Basing* will be added in Sect. 3.4) fits this schema for universal A-quantifiers, with an added argument for the subject (reflecting the fact that ‘know’ is a verb that takes a subject argument). Here is the schema for universal A-quantificational verbs that take a subject argument:

$$\llbracket A-Q \text{ verb} \rrbracket^c = \lambda E \lambda P \lambda x \lambda s. \text{ every } s' \text{ such that } L\text{— and } C(s') \text{ and } E(s'), \text{ is also such that } P(s')$$

Our proposed implementation instantiates this schema, with a new lexical restriction, L^{evid} , restricting to worlds consistent with the subject’s evidence (reflecting the distinctive meaning of ‘know’):

$$L^{evid} = \lambda x \lambda s \lambda s'. s' \text{ is consistent with } x\text{'s evidence in } s$$

Thus we view the evidential component of the semantic clause for ‘know’ in comparativist terms, as drawing a quantitative comparison between a contextually variable domain proposition and the scopal proposition said to be known:

$$(\text{First pass}) \llbracket \text{know} \rrbracket^c = \lambda E \lambda P \lambda x \lambda s. \text{ every } s' \text{ such that } L^{evid}(x, s, s') \text{ and } C(s') \text{ and } E(s'), \text{ is also such that } P(s')$$

Let’s correlate some of our epistemic and semantic concepts. On the epistemic side, the relevant alternatives theorist distinguishes between *the relevant options* (including the known proposition P) and *the relevant alternatives* (the relevant options in which P is false). The relevant options are encoded in *First pass* as the situations compatible with the contextual and explicit domain restrictions. Back on the epistemic side, the relevant alternatives theorist then requires that all of the relevant alternatives be eliminated. This requirement is encoded in *First pass* by the quantitative comparison, via the lexical requirement that the relevant options consistent with the subject’s evidence all be such that P .³⁶

³⁶ In Schaffer’s (2005a) contrastivist terminology, knowledge is a ternary relation of the form $KsPQ$, which may be read as ‘ s knows that P rather than Q .’ Contrastivism can be seen as the natural metaphysics of contextualist versions of *Relevant alternatives theory*, with the Q -slot holding the relevant alternatives to P . To know that P rather than Q is then to have evidence that eliminates Q , and to form a true belief that P on this basis. *First Pass* can be understood as implementing the evidential component of contrastivism, with the contrast proposition Q defined as follows: $Q = \lambda s.C(s) \text{ and } E(s) \text{ and } \sim P(s)$

First pass thus implements *Evidence*, understood in a contextualist way as per *Key*, which was the essential contextualist component of *Relevant alternatives theory* (Sect. 3.2). The implementation is orthodox, insofar as it fits the template of the A-quantifier ‘always’ (with minimal adjustments reflecting the fact that ‘know’ takes a subject argument, and imposes its distinctive lexical requirements). We still need to incorporate *Truth*, *Belief*, and *Basing*. But *First pass* forms the core of our solution to the semantic problem.

One might object to *First pass* on the grounds that, by treating *Evidence* in a quantificational manner, *First pass* incorporates the assumption that evidence is a merely *intensional*—as opposed to a *hyper-intensional*—notion. *First pass* treats evidence as being as fine-grained as the points of quantification (situations). It does not allow evidence to divide at any finer grain. Of course, just how coarse-grained this is depends on the metaphysics of situations: one can work with metaphysically or even mathematically impossible situations. But *First pass* does treat evidence as closed under (classical) entailment. If *P* entails *Q*, then satisfying the evidential requirement towards *P* entails satisfying the evidential requirement towards *Q*.³⁷

We think that evidence is indeed an intensional notion closed under entailment—or at least that the notion in play in our presupposed *Relevant alternatives theory* is best understood in this way. As to closure under entailment, suppose that one has evidence that eliminates all relevant alternatives to a given proposition *P*, and that *P* entails *Q*. Then it seems that one has all the evidence one could want for *Q*. Of course one might fail to *exploit* the evidence, but it is present all the same. From the perspective of *Relevant alternatives theory*, since *P* entails *Q*, then the relevant alternatives to *Q* must be a subset of the relevant alternatives to *P*. And since one has evidence that eliminates all the relevant alternatives to *P* (by hypothesis), one’s evidence thereby eliminates all the relevant alternatives to *Q*. *This is just what Evidence demands with respect to Q*. Again one might not exploit the evidence, but from the perspective of *Relevant alternatives theory* that can only be a failure of *Belief* or *Basing*, not a failure of *Evidence*.

That said we should acknowledge that there are reasons to prefer a hyper-intensional view of evidence. After all it still seems plausible that the logician can have evidence that a given theorem holds without having evidence that every theorem holds. We note that the very same issues arise for epistemic modals. It seems plausible that the logician can truly say something of the schematic form:

39. It might be false that _____

Where the blank is filled in by a (true) theorem that remains unproven, or whose proof has not yet reached her ears. So to whatever extent there is a problem of hyper-intensionality, it is if anything *confirmation* for our analogy between ‘know’ and modals that the same issue arises in both cases.

³⁷ *Proof:* Suppose that *x* satisfies the evidential requirement towards *P* in a context *c* where the domain proposition is *D*. Then every *D* situation consistent with *x*’s evidence is a *P* situation. By the supposition that *P* entails *Q*, every *P* situation is a *Q* situation. So every *D* situation consistent with *x*’s evidence is a *Q* situation, which is what it takes for *x* to satisfy the evidential requirement towards *Q* in *c* (*QED*).

There is a deep underlying problem surfacing, which is that there is not yet an accepted semantic framework for hyper-intensionality. This is a problem for everyone. We have no special solution on offer. So we proceed as best we can: we offer a semantics in the usual intensional idiom and trust that—should a framework for hyper-intensionality emerge—our semantics will prove smoothly extendible.

One might also object to *First Pass* in ways that will be familiar from standard epistemological objections to *Relevant alternatives theory*. For instance, one standard objection concerns the prospect of *cheap evidence* when a given proposition P has no relevant alternatives. Given that the evidential requirement on knowledge is the elimination of the relevant alternatives, this requirement will trivially be satisfied if there simply are no not- P situations, or if none happen to be relevant. We note the objection but will not attempt a reply: we are not trying to defend a relevant alternatives theory of knowledge, but only to implement it.³⁸

3.4 Comparativism for ‘knows’: incorporating truth, belief, and basing

So far our proposed contextualist implementation of *Evidence* (with the lexical restrictor unpacked) is:

$$(First\ pass) \llbracket know \rrbracket^c = \lambda E \lambda P \lambda x \lambda s. \text{ every } s' \text{ such that } s' \text{ is consistent with } x's \\ \text{evidence in } s, C(s'), \text{ and } E(s'), \text{ is also such that } P(s')$$

It remains to incorporate the *Truth*, *Belief*, and *Basing* components of *Relevant alternatives theory* into our denotation for ‘know.’ There are two main options: insert these inside the A-quantificational clause of *First pass*, or append these as a separate clause.

The first option of incorporating *Truth*, *Belief*, and *Basing* into the A-quantificational clause would yield a treatment of ‘know’ as a pure A-quantifier.³⁹ This would entail that knowledge is a merely intensional notion (as does any view on which attitude verbs are treated as verbal quantifiers: Sect. 2.4). Though (as discussed in Sect. 3.3) there is independent reason to think that some hyper-intensional generalization of A-quantification will be needed to handle epistemic modals. So we consider it premature to judge whether this is problematic.

This first option may be implemented in various ways, but here is an illustration for the sake of definiteness:

³⁸ Perhaps cheap evidence is not so worrisome. Cheap evidence is still not cheap knowledge, since the subject may not form the needed belief, or do so on a proper basis. (For a theory that does deliver cheap knowledge, see Lewis (1996, pp. 561–562); see Lihoreau (2008) for a line of objection to Lewis based on this fact.)

³⁹ This is in some respects the option that Lewis takes (modulo the distinction between D-quantification and A-quantification, and the fact that Lewis does not endorse *Belief* or *Basing* as stated). Essentially Lewis adds further lexical constraints on ‘know,’ some of which further contract the domain, but others of which ensure that the domain includes the situation (/world) of the subject (so as to capture *Truth*), and others of which ensure that the domain includes any situation (/world) the subject believes to obtain, or ought to believe obtains (so as to capture something in the neighborhood of *Belief* and *Basing*). Lewis thus upholds a purely quantificational semantics.

(*Second pass*) $\llbracket \text{know} \rrbracket^c = \lambda E \lambda P \lambda x \lambda s.$ every s' such that either (s' is consistent with x 's evidence in s , $C(s')$, and $E(s')$) or (s' is consistent with x 's properly based true beliefs in s), is also such that $P(s')$

Second pass implements *Truth*, *Belief*, and *Basing* by disjunctively expanding the domain to include situations consistent with x 's properly based true beliefs in s . The idea is to expand the domain so that if P is false, not believed, or not believed on a proper basis, then the domain should include a not- P situation, thereby falsifying the claim to knowledge.

But the treatment of *Belief* in *Second pass* is suspect, and the treatment of *Basing* in *Second Pass* is especially suspect. For the fact that x 's properly based beliefs entail P does not yet entail that x believes P , or that x believes P on that basis. Maybe there is a legitimate conception of belief on which the fact that x 's beliefs entail P entails that x believes P , but it is an even further stretch to imagine that there is any legitimate conception of basing—a notion intended to pick up on the actual mechanism of belief formation—that is blind to the actual mechanism of belief formation. We do not see how to fix this within a quantificational framework (leaving open whether a hyper-intensional generalization might in any way help).

The second option involves appending a separate clause for these additional factors. This second option may also be implemented in various ways, but here is an illustration for the sake of definiteness:

(*Third pass*) $\llbracket \text{know} \rrbracket^c = \lambda E \lambda P \lambda x \lambda s.$ every s' such that s' is consistent with x 's evidence in s , $C(s')$, and $E(s')$, is also such that $P(s')$; and x truly believes P on the basis of her evidence in s

Third pass implements *Truth*, *Belief*, and *Basing* by conjunctively adding the relevant requirements. As a result *Third pass* does not treat 'know' as an A-quantifier, but rather—in a novel way—as a hybrid with an A-quantificational aspect and a non-quantificational aspect. If one thinks that evidence is an intensional notion, but that belief is hyper-intensional and that the hyper-intensionality of knowledge comes from the hyper-intensionality of belief, then *Third pass* may be preferable. The A-quantificational aspect provides for context sensitivity, and the belief aspect provides for hyper-intensionality.

That said, we think—perhaps surprisingly—that neither *Second pass* nor *Third pass* implements *Truth*, *Belief*, or *Basing* in quite the right way. Neither fits the way 'know' interacts with explicit restrictors as in:

40. If it rains tomorrow I know my basement will flood

Obviously 40 can be true even if in fact my basement will not flood (it might not rain tomorrow). And likewise 40 can be true even if I do not believe that my basement will flood (I might not believe that it will rain tomorrow). This example shows that P (the proposition that my basement will flood) need neither be true nor believed. Instead P only needs to truly characterize the situations in which the explicit restrictive material obtains (the situations in which it does rain tomorrow), and all that needs to be believed on the basis of the evidence is that P truly characterizes these situations.

Accordingly our preferred implementation of *Truth*, *Belief*, and *Basing* goes as follows. For *Truth*, we stipulate (as a part of what we mean by ‘elimination’) that the subject’s evidence can never eliminate what is true, and add a semantic requirement that the contextual restrictor *C* be true. This ensures that the only restrictor argument which can be false in a true knowledge ascription is the explicit restrictor *E*. Since a true knowledge ascription—by *Evidence*—requires that *P* be true in all situations in which all three restrictor arguments are true, this entails that in a true knowledge ascription *P* is true given *E*. Where there are no explicit restrictive clauses, *E* is the trivial proposition true in all situations (Sect. 3.1), so where there are no explicit restrictive clauses we recover the usual idea that *P* must be true. But where there are explicit restrictive clauses we allow that *P* may be false as long as it is *true given E*, just as 40 seems to demand.

Turning to *Belief* and *Basing*, our idea is that the doxastic side of the ledger should mirror the evidential side of the ledger. What *Evidence* demands on the evidential side is that, presupposing *C* and *E*, the subject’s evidence provides (conclusive) evidence against any remaining prospect of falsity for *P*. So on the doxastic side of the ledger the subject should be required to have (conclusive) belief in *P* presupposing *C* and *E*. In other words, what we require for *Belief* is that the subject have (conclusive) belief in *P given C and E*, and what we require for *Basing* is that this belief be properly based on the evidence. These requirements may be understood as a state of properly based *restricted certainty*: one must be disposed to certainty on the basis of the evidence that, *given that one of the relevant options holds*, *P holds* (Schaffer 2005a, pp. 255–256).

Putting this together (and continuing with the idea of separate clauses as seen in *Third pass*) yields:

(*Final pass*) $\llbracket \text{know} \rrbracket^c = \lambda E \lambda P \lambda x \lambda s. \text{ every } s' \text{ such that } s' \text{ is consistent with } x\text{'s evidence in } s, C(s'), \text{ and } E(s'), \text{ is also such that } P(s'); \text{ and } C(s); \text{ and } x \text{ believes } (P \text{ given } C \text{ and } E) \text{ on the basis of her evidence in } s^{40,41}$

For the sake of definiteness we will work with *Final pass* for what remains, but would just reiterate that most of the details found in *Final pass* are inessential to our purposes. Our core claim is just that ‘know’ has an A-quantificational aspect, which involves a quantitative comparison between two propositions (the first of which is expressed in a contextually variable way, connected to the question under discussion). *First pass* is just a fairly orthodox (albeit partly simplified) way to

⁴⁰ *Final pass* takes advantage of our distinct *C* and *E* variables (Sect. 3.1). Had we lumped *C* and *E* together into a single conjoint restrictor argument, ‘know’ would presumably only have access to this conjunction, and would not be able to “see” what restriction was due to relevance and what was due to explicit restriction.

⁴¹ Recall that *Final pass* is built on a simplified version of Amherst semantics that does not incorporate parthood or minimality (Sect. 3.1). Incorporating these complications yields:

(*Final pass**) $\llbracket \text{know} \rrbracket^c = \lambda E \lambda P \lambda x \lambda s. \text{ Every } s' \text{ such that } s' \text{ is a minimal situation consistent with } x\text{'s evidence in } s, C(s'), \text{ and } E(s'), \text{ is part of a situation } s'' \text{ such that } s'' \text{ is consistent with } x\text{'s evidence in } s, C(s''), E(s''), \text{ and } P(s''); \text{ and } C(s); \text{ and } x \text{ believes } (P \text{ given } C \text{ and } E) \text{ on the basis of her evidence in } s.$

implement our core claim, and *Final pass* merely represents our preferred way to complete the matter. We can perhaps best capture the contextualist core of our semantics with the following schema:

(*Final schema*) $\llbracket \text{know} \rrbracket^c = \lambda E \lambda P \lambda x \lambda s. \text{ every } s' \text{ such that } s' \text{ is consistent with } x\text{'s evidence in } s, C(s'), \text{ and } E(s'), \text{ is also such that } P(s'); \text{ and } \underline{\hspace{1cm}}$

Final schema is just *First pass* with a blank space added for any additional requirements to be conjoined. Equally, *Final schema* is just *Final pass* with a blank space in place of its *Truth*, *Belief*, and *Basing* conjuncts.

4 Arguments for epistemic comparativism

We have provided a comparativist semantics for ‘know’ in the form of *Final pass*. To repeat:

(*Final pass*) $\llbracket \text{know} \rrbracket^c = \lambda E \lambda P \lambda x \lambda s. \text{ every } s' \text{ such that } s' \text{ is consistent with } x\text{'s evidence in } s, C(s'), \text{ and } E(s'), \text{ is also such that } P(s'); \text{ and } C(s); \text{ and } x \text{ believes } (P \text{ given } C \text{ and } E) \text{ on the basis of her evidence in } s$

It remains to discuss whether *Final pass*—or anything fitting *Final schema*, or even anything extending *First pass*—is viable. If so the semantic problem for *Epistemic contextualism* is resolved.

The analogy with A-quantifiers which inspired our semantics will provide further help in assessing our semantics. A-quantifiers have certain characteristic features, and we will now argue that ‘know’ displays these features. We will then return to our five desiderata for a plausible contextualist semantics (Sect. 2.4), to show how our semantics satisfies every desideratum.

4.1 Question sensitivity

Our primary argument for comparativism has already surfaced (Sects. 1.2, 2.4): ‘know’ and A-quantifiers display a parallel sort of question sensitivity. Let’s try to get underneath how this works. To begin with, we take it that context includes a parameter for *the question under discussion*. This parameter helps explain a wide range of phenomena. For instance, Ginzburg (1996) discusses an actual conversation that starts with the question of how old a given person is, and ends many turns later with the elliptical utterance: “Seventy two.” As Ginzburg (1996, p. 415) explains, the question must be available to interpret the distant ellipsis:

[S]ince the discussion of a single question can last over several turns, and elliptical contributions are possible, in principle, arbitrarily far away from the turn in which the question was posed, what will be needed is a notion of context which can express the fact that a particular question is (still) under discussion,...

The question under discussion parameter has also been used to explain and clarify the Gricean maxims, rules of turn taking, and discourse coherence.⁴² Thus Carlson (1983; cf. Roberts 2004, p. 209) speaks of the discourse game as built around questions (setup moves) and answers (payoff moves), and Roberts (2004, p. 216) explains and clarifies the Gricean maxim of *Relevance* in terms of addressing the question:

A move *m* is RELEVANT to the question under discussion *q* iff *m* either introduces a partial answer to *q* (*m* is an assertion) or is part of a strategy to answer *q* (*m* is a question subordinate to *q* or an imperative whose realization would plausibly help to answer *q*.)

Now we add the hypothesis that the contextual restriction on A-quantifiers is constrained by the question under discussion (von Fintel 1994, 2004). Thus imagine that Claire and Elaine are rival jewel thieves, that Claire only takes diamonds, but that Elaine takes both diamonds and rubies. In the situation at hand Claire has stolen the diamonds. Now consider the question of what Claire steals. In such a context:

37. Claire always steals the diamonds

seems true, since diamonds are what Claire always steals. But now consider the question of who stole the diamonds, and reconsider 37 in such a context. In such a context 37 seems false, since sometimes it is Elaine who steals the diamonds. So it seems that the question under discussion is constraining the implicit situation domain. (*WhoAlways* and *WhatAlways* in Sect. 2.4 make this very same point.)

The hypothesis that the contextual restriction on A-quantifiers is constrained by the question under discussion allows other factors to play a role. But the simplest hypothesis—which we hereby adopt for definiteness—is that the contextual restriction just *is* the presupposition of the question under discussion. That is, the question under discussion in a given context denotes a set of possible answers, and the *C-restrictor* argument is the proposition true of all and only the situations some possible answer to the question is true of.⁴³ Thus, if the question under discussion is the question of who stole the diamonds, the contextual restriction on the domain is the proposition that someone (among the contextually relevant individuals) stole the

⁴² Some further applications: Beaver and Clark (2008) use the question under discussion parameter to evaluate focus variables (Sect. 4.2), Schoubye (2009) uses it to predict when existential presupposition failure triggers falsity intuitions, Potts (2011) uses it to explain the felt “negativity” of negation, Simons et al. (2010) use it to explain implication projection, and Schaffer (2011, Sect. 4) uses it to explain disagreement dynamics over taste and epistemic modality.

⁴³ We assume that a question like ‘Who stole the diamonds?’ presupposes that someone (among the contextually relevant individuals) stole the diamonds. In support of this we note the felt contrast between ‘Who stole the diamonds?’ and ‘Who, if anyone, stole the diamonds?’ Horn (1972) calls clauses such as ‘if anyone’ *suspenders* and argues that they should be understood as suspending presuppositions (cf. AnderBois 2009). The answer set view of questions we are operating with, which traces back to Hamblin (1973), treats questions as having presuppositions. The partition view of questions, pioneered by Groenendijk and Stokhof (1984), does not, although one can work with a partition view that does not partition the entirety of logical space but only some contextually selected subset (e.g. the common ground).

diamonds; while if the question under discussion is the question of what Claire stole, the contextual restriction of the domain is the proposition that Claire stole something (among the contextually relevant items).

What we have at this point is an account of the mechanism of context sensitivity in play with A-quantifiers. The account posits a certain type of semantic argument for A-quantifiers—the *C-restrictor* variable (Sect. 3.1)—evaluated in a contextually variable way as the presupposition of the question under discussion. We have already argued that ‘know’ features a similar question sensitivity. Thus recall:

(*Who*) Claire has stolen the diamonds. Ann and Ben are wondering who stole the diamonds, and Ann finds Claire’s fingerprints all over the safe. So Ann says to Ben:

1. I know that Claire stole the diamonds

(*What*) Claire has stolen the diamonds. Ann and Ben are wondering what Claire stole, and Ann finds Claire’s fingerprints all over the safe. So Ann says to Ben:

1. I know that Claire stole the diamonds

And so the most natural hypothesis is that the same mechanism is at work: ‘know’ also features a *C-restrictor* variable, evaluated in a contextually variable way as the presupposition of the question under discussion, and playing a domain restricting role. Our semantics thus explains the question sensitivity of knowledge ascriptions, and the parallel between ‘know’ and ‘always’ in this respect.

Note that in *Final pass* we implemented *Truth* by requiring that *C* be true (Sect. 3.4). Given our hypothesis that *C* is the presupposition of the question under discussion, this requirement becomes the requirement that the question under discussion does not have a false presupposition. To test whether this is plausible, imagine that the diamonds were not stolen at all (perhaps the owner hid them to collect the insurance money), but that Ann and Ben (being misled) are wondering who stole the diamonds. We predict that, no matter how much fingerprint or other evidence Ann might have, it would still be false for Ann to say:

1. I know that Claire stole the diamonds

We consider this a plausible result.

4.2 Association with focus

Our second argument for comparativism concerns *association with focus*. By ‘focus’ we mean a certain informational prominence, typically realized in English by nuclear pitch accent.⁴⁴ Focus plays various linguistic roles including a pragmatic

⁴⁴ Other languages realize focus in different ways. For instance, Hungarian realizes focus by a distinguished syntactic position (immediately preceding the verb), and Navajo realizes focus by a distinguished morpheme (‘-ga’).

role with respect to felicitous answers. Thus suppose that the question has arisen as to who stole the diamonds. Then:

41. CLAIRE stole the diamonds

is a felicitous answer but:

42. Claire stole THE DIAMONDS

is not. Whereas if the question has arisen as to what Claire stole then 42 is felicitous but 41 is not. In general, felicitous focus marking in an answer to a *wh*-question must exhibit *congruence*, in that nothing other than the constituent corresponding to the *wh*-word in the question can be in focus within the answer. This much is merely pragmatic: the non-congruent answers are infelicitous but still true.⁴⁵

Focus plays another role—this time semantic—in connection with a special class of expressions that exhibit “association with focus,” in the scope of which focus differences can make for truth-conditional differences. A classic example of an expression in this special class is ‘only.’ Thus compare:

43. Claire only stole THE DIAMONDS in Paris

44. Claire only stole the diamonds IN PARIS

43 is true if and only if Claire stole the diamonds and nothing else in Paris, but 44 is true if and only if she stole the diamonds in Paris and nowhere else. And so if the sum total of all of Claire’s thefts were the diamonds in Paris and the diamonds in Shanghai, then 43 would be true but 44 false.

A viable account of focus must explain the linguistic roles focus plays, including both congruence and association. We take it that the discourse function of focus is to identify the question under discussion at the time of the utterance (von Stechow 1994; Roberts 1996, 2004; Beaver and Clark 2008). The algorithm is simple: to find the logical form of the congruent question, just replace the focused constituent with an appropriate *wh*-word. To illustrate, the discourse function of the focus in 41 is to identify the question under discussion as the question of who stole the diamonds. This is a function worth serving, since the question under discussion plays such crucial roles (Sect. 4.1), but is often implicit, liable to shift, and subject to discourse negotiation.

The explanation of congruence is immediate. To illustrate, if the question is who stole the diamonds, then 41 is a felicitous answer but 42 is not. This is because 41 correctly identifies the question under discussion while 42 does not. The explanation of association follows naturally, given a semantic treatment of questions as sets of possible answers (Sect. 4.1), and a semantic treatment of the special class of expressions that associate with focus as sensitive to the set of possible answers by taking a semantic argument anaphoric on this set. For example, we could say that ‘only ϕ ’ is true just in case ϕ expresses the sole (or perhaps the strongest) true answer

⁴⁵ Non-congruent answers are true in English, but may be false in other languages—such as Hungarian—with focus constructions that allow for a truth-conditional effect even without any expression that associates with focus (Kiss 1998).

to the question under discussion (Beaver and Clark 2008, chap. 10). Such an account of ‘only’ explains association with focus as a byproduct of question sensitivity.

Returning to A-quantifiers, these turn out to be in the special class of expressions that associate with focus (Rooth 1996, p. 272), which is just as we should expect given their question sensitivity. Thus suppose that Claire and Elaine are rival diamond thieves, both of whom only take diamonds. Then:

45. Claire always steals THE DIAMONDS

seems true, but:

46. CLAIRE always steals the diamonds

seems false. Given that A-quantifiers include a context sensitive *C-restrictor* variable anaphoric on the presupposition of the question under discussion (Sect. 4.1), we can derive the difference between 45 and 46. With 45, assuming felicity (a charitable default assumption), the question under discussion must concern what Claire steals. So our Amherst semantics will require that every actual situation in which Claire steals something is one in which she steals the diamonds. The requirement is met and so 45 is true. (Whereas if we imagined instead that Claire sometimes steals rubies, then 45 would be false.)

With 46, assuming felicity, the question under discussion must concern who steals the diamonds. So our Amherst semantics will require that every actual situation in which someone steals the diamonds is one in which Claire steals the diamonds. Since Elaine sometimes steals diamonds, 46 is false. (Whereas if we imagined instead that Elaine only steals rubies, and that Claire and Elaine are the only salient characters, then 46 would be true.) Though notice that we assumed felicity to derive the falsity of 46. But if we imagine someone actually saying 46, then the assumption that they were speaking felicitously would lead to the conclusion that they were speaking falsely, and so charity might ultimately lead us to suspend judgment as to whether they were speaking felicitously but falsely, or infelicitously but truly. So we may end up somewhat unsure about how to evaluate an actual utterance of 46. This actually accords with our intuitions when we try to imagine someone saying 46, though we acknowledge that intuitions may be delicate. In any case we take it that our semantics predicts a *relevant difference* between 45 and 46.

We have already—if only in passing—suggested that ‘know’ is also in the special class of expressions that associate with focus, by providing empirical data (Sect. 1.2) on the closely related phenomena of cleft constructions. To this end recall that people, when given the vignette *New*, seemed happy to assent to 4 and then immediately deny 5, and equally happy to deny 5 and then immediately assent to 4:

4. Mary knows that it was Peter that stole rubies
5. Mary knows that it was rubies that Peter stole

Moving from cleft constructions to focus:

(*WhoFocus*) Claire has stolen the diamonds. Ann finds Claire’s fingerprints all over the safe. So Ann says to Ben:

47. I know that CLAIRE stole the diamonds

(*WhatFocus*) Claire has stolen the diamonds. Ann finds Claire's fingerprints all over the safe. So Ann says to Ben:

48. I know that Claire stole THE DIAMONDS

We think that 47 seems true in *WhoFocus*, but that 48 seems false in *WhatFocus*.⁴⁶ Given that 'know' has the A-quantificational aspect we posit, we can derive the difference between 47 and 48 in a way parallel to the derivation of the difference between 45 and 46.

Thus with 47, assuming felicity (a charitable default assumption), Ann and Ben must have been wondering about who stole the diamonds. This reduces the status of 47 in *WhoFocus* to the status of 1 in *Who*. Our semantics predicts that the proposition expressed by Ann is (roughly speaking) that given that one of the contextually salient suspects stole the diamonds, Ann's evidence entails that Claire stole the diamonds. In light of what we are told about the case this seems true.

With 48, assuming felicity, Ann and Ben must have been wondering about what Claire stole. This reduces the status of 48 in *WhatFocus* to the status of 1 in *What*. Our semantics predicts that the proposition expressed by Ann is (roughly speaking) that given that Claire stole one of the contextually salient items, Ann's evidence entails that Claire stole the diamonds. This should seem false. Though again—as with 46—we actually predict something slightly more subtle. Whereas the question under discussion is explicit with 1 in *What*, with 48 in *WhoFocus* we need an assumption of felicity. Since this leads to an imputation of falsehood, we may end up somewhat unsure about the case overall (a felicitous falsehood, or an infelicitous truth?) This accords with our intuitions: 1 indeed seems more clearly false than 48. But again we acknowledge these intuitions to be delicate, and think the most important thing is for the semantics to predict a relevant difference between 47 and 48.

We would add that association with focus is only predicted by a contextualist view, and moreover—given the connection between focus and the question under discussion—only predicted by a contextualist view that accords the question under discussion a semantic role (which for us goes via the evaluation of the *C-restrictor* argument). Association with focus just is an indirect effect of question sensitivity.

⁴⁶ The claim that 'knows' associates with focus traces back to Dretske (1981, p. 373; cf. Sanford 1991), who gives the following illustration: "Someone claiming to know that Clyde *sold* his typewriter to Alex is not (necessarily) claiming the same thing as one who claims to know that Clyde sold his typewriter to Alex... A person who knows that Clyde *sold* his typewriter to Alex must be able to rule out the possibility that he *gave* it to him, or that he *loaned* it to him... But he needs only a nominal justification, if he needs any justification at all, for thinking it was Alex to whom he sold it." Drawing on Dretske, von Stechow (1982, p. 3) then claims: "Interrogative-embedding operators are *structure-sensitive* ('focus-sensitive'): 'God knows whether Ede loves SENTA' will have different truth-conditions from 'God knows whether EDE loves Senta'."

4.3 Scoping over restrictive clauses

Our third argument for comparativism concerns *scoping over restrictive clauses*. What we have in mind has already surfaced in Sect. 3.1 with:

38. Claire always steals the diamonds when she is bored

In 38 we treated the ‘when’-phrase as contributing to the value of the *E-restrictor* argument of ‘always’ (we evaluate *E-restrictor* arguments as the proposition expressed by the conjunction of any explicit restrictive clauses plus the trivial proposition). A-quantifiers like ‘always’ interact in distinctive ways with such restrictive clauses, tending to scope over them even when surface form suggests otherwise. Thus, surface form notwithstanding, 38 seems equivalent to:

49. When Claire is bored she always steals the diamonds

Or imagine that Ann and Ben are playing poker:

50. If Ann holds three of a kind she usually wins

51. If Ann holds the better hand she must win

In both 50 and 51 there is a strong preference for taking the relevant A-quantifier to scope over the ‘if’-clause, surface form notwithstanding. For instance, 50 is most naturally read as saying that usually, in those situation in which Ann holds three of a kind, she wins in those situations. Indeed scoping behavior constitutes one of the main types of evidence for imputing tripartite quantifier-restrictor-scope structure to A-quantifiers: the best explanation for the natural scopings in 49–51 is that ‘always,’ ‘usually,’ and ‘must’ are quantifiers which ‘if’ and ‘when’ clauses serve to restrict.

Of course the tendency for the A-quantifier to take wide scope over the restrictive clause is defeasible. In this vein consider:

52. If Claire steals the diamonds she always steals the diamonds⁴⁷

52 seems most naturally to generate a narrow scope reading (perhaps the current situation is being treated as a test for Claire’s tendencies, and it is being asserted that if Claire steals the diamonds in this situation, she does so in all situations). For a wide scope reading would involve the assertion of a very simple tautology (every situation in which Claire steals the diamonds is part of a situation in which she steals the diamonds), which is pragmatically discouraged.⁴⁸ So we would sum up the distinctive way in which A-quantifiers interact with restrictive clauses as follows: A-quantifiers take wide scope over restrictive clauses when plausibility permits.

⁴⁷ Modal versions of these examples are discussed in Frank (1996) and Zvolenszky (2002).

⁴⁸ Kratzer (2012, revising her 1991) expresses some doubt as to whether the unavailability of wide scope readings for modals can always be explained pragmatically, citing cases like ‘If he has a kitchen, he can cook.’ It is true that such a sentence is not naturally read as saying of someone that he has a cooking ability conditional on owning a kitchen. But this would be a rather bizarre ability. Having the ability to cook on the condition of finding a kitchen is a more usual condition and, as expected, we do get the wide scope reading for ‘If he finds a kitchen, he can cook.’

‘Know’ patterns with A-quantifiers in taking wide scope over restrictive clauses when plausibility permits. Thus imagine that Ann and Ben are playing poker, and that Ann holds a pair of kings:

53. If Ben is holding a pair of queens Ann knows she will win

The most natural interpretation of 53 has ‘know’ taking wide scope over the ‘if’-clause, surface form notwithstanding. This is the reading on which Ann is being said to know that she will win in case Ben holds a pair of queens. Indeed the interpretation on which ‘know’ takes narrow scope under the ‘if’-clause is bizarre, making Ben’s hidden hand some sort of condition for Ann’s knowledge that she will win.

Likewise imagine that, while Ann holds her pair of kings, Ben is awaiting his final cards:

54. If Ben will receive a pair of queens Ann knows she will win

Again the most natural interpretation of 54 has ‘know’ taking wide scope over the ‘if’-clause—Ann is being said to know that she will win in case Ben will receive a pair of queens. The interpretation on which ‘know’ takes narrow scope under the ‘if’-clause takes wide scope is bizarre, making’s Ben’s receiving a pair of queens in the future a condition on Ann’s present knowledge.

Or consider the following story:

(*KnowCards*) Ann and Ben are playing poker, and Frank is watching. Frank circles behind Ann and sees that she holds a pair of kings. He thinks: “If Ben is holding a pair of queens Ann knows she will win.” So he circles behind Ben and sees that Ben is holding a pair of queens. So he thinks:

55. Ann doesn’t yet know that she will win, but she will win

Frank’s thoughts seem perfectly coherent in *KnowCards*, but notice that Frank has thought:

- If Ben is holding a pair of queens Ann knows she will win
- Ben is holding a pair of queens
- Ann does not yet know that she will win

If the first of these thoughts is interpreted with narrow scope ‘know’ (as per surface form) then Frank’s thoughts would blatantly violate *modus ponens* reasoning. The reason why Frank’s thoughts are coherent is that the first of his thoughts is naturally read with wide scope ‘know.’ All Ann knows is that she will win in case Ben holds a pair of queens. Since Ann does not yet know that this is the case, she does not yet know that she will win.

Of course we are only claiming that ‘know’ takes wide scope over restrictive clauses *when plausibility permits*. For a case in which plausibility does not permit such scoping, consider:

56. If Ann holds a pair of kings she knows she holds a pair of kings

(Think of 56 said as a way of conveying that Ann knows the rules of poker.) On the reading on which ‘know’ takes wide scope, Ann is merely being said to know the tautology that she holds a pair of kings in case she holds a pair of kings. That would be a strange attribution, at least in any ordinary context.

We remain neutral on whether adverbial quantifiers, modals, and ‘know’ exhibit a tendency to the same degree for taking wide scope over restrictive clauses. Perhaps there are differences (if so we leave open what might explain them). Our core point is that wide scope readings are *possible*.

It has been known since Lewis (1975) that no plausible connective predicts the right truth-conditions for adverbially quantified conditionals where the quantifier takes wide scope. This is a crucial datum that motivates Lewis and Kratzer to treat ‘if’-clauses uniformly as explicit restrictors on quantificational domains. Our semantics follows this orthodox line, and can thereby handle conditionals within the scope of ‘always’ as well as ‘know.’ Those who would resist treating ‘know’ as having a quantificational aspect owe some other account of the most natural readings of 53–55.⁴⁹

4.4 Domain coordination

Let’s look at one last feature of A-quantifiers, which is a capacity to participate in *domain coordination*. By ‘domain coordination’ we mean cases in which a second A-quantifier inherits its domain from a first. (This is akin to *binding* for A-quantificational restrictor arguments.) To illustrate what we have in mind, consider:

57. Claire always steals the diamonds and she never leaves fingerprints

The second conjunct of 57 is most naturally interpreted as quantifying over the set of situations invoked in the first: if the first conjunct says that whenever Claire steals something she steals the diamonds, the second then says that whenever she steals something she doesn’t leave fingerprints; if the first conjunct says that whenever someone steals the diamonds Claire steals them, the second then says that whenever someone steals the diamonds Claire does not leave fingerprints. Indeed, if Claire is a careless thief who often leaves fingerprints *when thieving*, and merely never leaves fingerprints *when washing the dishes*, then 57 should sound false.

Or imagine that Claire usually targets diamonds and is usually very careful, but once stole the rubies and left fingerprints:

58. Claire often steals the diamonds and she never leaves fingerprints

58 should sound false, since Claire does sometimes leave fingerprints. This requires that the second conjunct of 58 be interpreted as quantifying over all the situations in which Claire steals anything, and not merely those in which she steals the diamonds.

⁴⁹ Gillies (2010) offers a dynamic semantics that can handle conditionals within the scope of epistemic modals, but we see no straightforward way to generalize his proposal to adverbial quantifiers or ‘know.’ For a discussion of the difficulties Gillies’s account faces with adverbial quantifiers, see Khoo (2011).

The *C-restrictor* variable allows for a natural and simple explanation for domain coordination. Given that the question under discussion does not shift, and given our hypothesis that the *C-restrictor* is anaphoric on the question under discussion, we derive a common domain.⁵⁰

We will now argue that ‘know’ can also participate in relations of domain coordination, both with other occurrences of ‘know’ and with other A-quantifiers like ‘must’ and ‘always.’ To begin with, compare:

1. I know that Claire stole the diamonds
59. I know that Claire must have stolen the diamonds

1 and 59 seem—on the most natural readings—to be truth-conditionally equivalent (Stephenson 2007, Yalcin 2007). Of course 59 suggests that Ann’s evidence is somehow indirect, perhaps involving reasoning or testimony, while 1 allows that Ann might have direct perceptual evidence. But the suggestion of evidential indirectness is cancelable and so presumably not truth-conditional:

60. Ann knows that Claire must have stolen the diamonds, in fact she saw her do it

The puzzle is to explain the truth-conditional equivalence of 1 and 59 (on their most natural readings). It is not puzzling why 59 should entail 1, but quite puzzling how 1 could entail the seemingly stronger 59.

If the ‘must’ in 59 can inherit its domain from ‘know,’ then we can explain the entailment from 1 to 59. Given 1 and our semantics, we have it that all the relevant options compatible with the speaker’s evidence are situations in which Claire stole the diamonds. Moreover our semantics also entails that, if all the relevant options compatible with the speaker’s evidence are situations in which Claire stole the diamonds, then all the relevant options compatible with the speaker’s evidence are situations in which all the relevant options compatible with the speaker’s evidence are situations in which Claire stole the diamonds. Given domain coordination, that will be just what 59 requires, at least with respect to the evidential aspect of knowledge encoded in *First pass*. So assuming that any extra requirements like *Belief* are also satisfied, 59 will hold.⁵¹

Leaving epistemic modals behind, we observe domain coordination between ‘always’ and ‘know’ in sentences such as:

⁵⁰ We also take advantage of our hypothesis that ‘always’ and ‘never’ share a lexical restriction to actual situations. In coordination between A-quantifiers with different lexical restrictions we could either hold that the domain coordination is merely partial, or posit that the second *C-restrictor* argument is co-indexed with the first. This second option takes further advantage of the tripartite restrictor structure we posit (Sect. 3.1), by allowing coordination of arguments inside this structure.

⁵¹ It is possible for a subject to believe the scopal material in 1 but not 59 (perhaps the subject has heterodox beliefs about modals) or to believe 59 but for the wrong reasons (perhaps the subject doesn’t notice a connection to 1 but believes 59 on the basis of her horoscope). So the sort of entailment seen from 1 to 59 is best understood as only valid given certain background assumptions about the subject’s belief formation (akin to the way closure inferences work). This is as our semantics predicts.

61. Claire always steals the diamonds and Ann knows that Claire always steals the diamonds

In a context in which ‘always’ quantifies over situations in which someone has stolen the diamonds, Ann needs evidence as to who stole the diamonds (e.g. fingerprints). But in a context in which ‘always’ quantifies over situations in which Claire has stolen something, Ann needs evidence as to what Claire stole.

We also observe domain coordination between multiple occurrences of ‘know’ in:

62. Ann knows that Claire stole the diamonds, and Ben knows that too.

Essentially Ann and Ben will be required to rule out the same range of relevant alternatives. If Ann has only the fingerprint evidence, and Ben has only the evidence that diamonds were stolen, then there will be contexts in which the first conjunct of 62 is true, and contexts in which the second conjunct of 62 is true, but no natural context in which the conjunction is true. Indeed, domain coordination between multiple occurrences of ‘know’ is already seen in the abominability of conjunctions like:

12. Ann does not know that she is not a bodiless and therefore handleless brain-in-a-vat, but she does know that she has hands

It is because the second occurrence of ‘know’ naturally inherits the domain of the first that we do not find a natural felicitous reading on which the relevant alternatives contract in mid-sentence.

Our comparativist semantics for ‘know’ explains how domain coordination is possible. Indeed we do not see how to explain the coordinated readings of 59–62 (inter alia) save through a semantics for ‘know’ that posits domain arguments that might coordinate.

Putting together the four arguments of Sects. 4.1–4.4, we claim:

- Comparativism explains the question-sensitivity of ‘know’
- Comparativism explains the association of ‘know’ with focus
- Comparativism explains the interaction of ‘know’ with restrictive clauses
- Comparativism explains domain coordination for ‘know’

Pending alternative explanations, we conclude that comparativism provides the best explanation for these four phenomena. ‘Know’ looks to bear the distinctive marks of A-quantification.

4.5 Desiderata for plausibility met

It remains to show—as our final argument for comparativism—that we have met all the desiderata for a plausible semantics outlined in Sect. 2.4. Recall that the desiderata were:

- A plausible contextualist semantics should not predict smooth tracking across contexts
- A plausible contextualist semantics should predict overall discourse constraints that preclude free shifting within a discourse

- A plausible contextualist semantics should not predict smooth interaction with comparative and degree morphology, and should not allow paraphrases up and down an associated scale
- A plausible contextualist semantics should not require a typically overt expression to constrain the context sensitivity
- A plausible contextualist semantics should predict question sensitivity

Satisfaction of the third desideratum is immediate: we do not posit a degree argument or anything else that would support smooth interaction with comparative and degree morphology, or semantic interpretation relative to an underlying scale. Satisfaction of the fourth desideratum is immediate from our characterization of A-quantifiers in Sect. 2.4, and satisfaction of the fifth desideratum has been argued for in Sect. 4.1. So what remains to be shown is satisfaction of the first two desiderata.

Starting with the first desideratum, it is widely thought that we do find “semantic blindness” with certain standard A-quantifiers, namely *modals* (Sect. 3.2). Indeed, as relativists such as MacFarlane (2011) and Egan et al. (2005) have detailed, epistemic modals exhibit all the tracking problems across contexts. For instance, we do not naturally adjust epistemic modals in indirect speech reports. If Ann says:

63. It might rain

this can be reported quite generally as:

64. Ann said that it might rain

even if the context has shifted. There is also felt cross-contextual disagreement with epistemic modals. If Ann says 63, and the meteorologist who is certain that it will not rain overhears Ann, the meteorologist may well regard what Ann said as false. Given our orthodox assumption that epistemic modals are A-quantifiers, we thereby have ample precedent to posit semantic blindness.⁵²

We think the above point about epistemic modals is sufficient for our purposes, but would add that adverbial quantifiers seem to exhibit much the same pattern. If Ann says:

36. Claire always steals the diamonds

this can be homophonically reported:

65. Ann said that Claire always steals the diamonds

⁵² We would emphasize that the charges of semantic blindness for both ‘know’ and for modals are based on armchair assessments of what is intuitive, and are thus open to empirical reconsideration. Indeed, Knobe (personal communication), together with Yalcin, conducted an experiment in which people were asked to evaluate an expert who says (in response to misleading evidence planted by a mobster who has faked his own death to evade the police): ‘Fat Tony might be dead.’ People overwhelmingly rejected the claim that what this expert said was false, but were all over the map as to whether what this expert said was true. Perhaps the armchair assessments of semantic blindness for modals have been overstated, and the same may be true for ‘know.’ Unfortunately we lack the analogous empirical data on knowledge ascriptions to discuss matters further.

even if the context has shifted.⁵³ Once one has homophonic indirect reports, full semantic blindness looks to come along for the ride. For if we can report Ann as having said that Claire always steals the diamonds in any context, including one in which 36 expresses a falsehood, then we can report Ann via 65 as having said something false. So if Ann says 36 and Ben overhears Ann while considering some situations that are not parts of situations in which Claire steals the diamonds, Ben may well likely regard what Ann said as false.

Of course one might take this pattern of semantic blindness across the A-quantifiers as reason to reject our assumed contextualism for A-quantifiers. But we think that this would be a mistake. Contextualism for A-quantifiers makes no predictions one way or another about our liabilities to error. The permissibility of homophonic indirect reports is just *data to be explained*. One should not a priori assume immunity to error with this (or any other) form of contextual variability. Error-theoretic explanations are legitimate explanations, if they can be empirically sustained. There is good independent evidence to think that epistemic modals and adverbial quantifiers have a common A-quantificational structure, exhibit a common form of contextual variability (sensitivity to the question under discussion), and generate a characteristic sort of liability to cross-contextual error. Knowledge ascriptions display the same pattern of contextual variability. If they also display the same pattern of error, then this is really just further evidence that ‘know’ has an A-quantificational aspect. So we think that semantic blindness should if anything count as a fifth mark of A-quantification.

Given that semantic blindness is so often taken as a “killer” objection to *Epistemic contextualism*, we think it may be worth distinguishing a bad from a good form of the objection. The bad form of the objection takes the premise that any contextual sensitivity in the language should be *obvious*. The underlying idea is semantic transparency: if linguistic behavior suggests that we are unaware of some putative feature of meaning then we should conclude that the feature does not exist. But semantic transparency is implausible. Virtually every sophisticated semantic theory posits all sorts of non-transparent features. Non-obvious context sensitivity is just more of the same. (Indeed we suspect that those who endorse the transparency premise must ultimately be the sort of radical invariantists who only allow for context sensitivity with core indexicals and demonstratives.)

The good form of the semantic blindness objection does not begin with any (implausible) assumption of semantic transparency or immunity from error, but rather invokes empirical arguments that we are not liable to certain sorts of errors with certain sorts of expressions. Indeed (as noted in Sect. 2.1) most invariantists work with the good form of the objection, and give plausible arguments that we do not make certain sorts of cross-contextual errors with indexicals, from which they

⁵³ Cappelen and Lepore (2005, pp. 94–96) hold that all utterances—save those with core indexicals and demonstratives—can be truthfully reported via disquotation in any context whatsoever. Given moderate contextualism in semantics (as we are presupposing: Sect. 1.3), it follows that either Cappelen and Lepore are wrong about homophonic reporting practices, or that semantic blindness is commonplace. For the record we actually suspect that both are the case. We suspect that homophonic indirect reporting fails for D-quantifiers (as we suggested in Sect. 2.3), but is available for A-quantifiers.

conclude that ‘know’ is not an indexical. *We agree with this argument.* We agree that semantic blindness undermines the analogy with indexicals. The anti-contextualist who would run the good form of the semantic blindness objection against our semantics must provide empirical arguments that we are not liable to the cross-contextual errors at issue with paradigmatic A-quantifiers. Our point is that we *do* actually seem to make this sort of error with certain paradigmatic A-quantifiers.⁵⁴

Turning to the second desideratum concerning discourse-level constraints, adverbial quantifiers are subject to discourse-level constraints since they take their contextual restrictors from the discourse-structuring question under discussion. This is in contrast to D-quantifiers which have overt restrictors that do not play any direct role in discourse structuring. For instance, it would sound terrible for Ann to say:

66. Claire never steals anything, and she always steals the diamonds

even if it is true that Claire never steals anything when dividing the loot with her accomplices, and always steals the diamonds when robbing jewelry stores. (Our point with 66 is really just a further example of our point with 58 about the strong preference for interpretations with domain coordination). To see how A-quantifiers work differently from D-quantifiers in this respect, compare Stanley’s (2005, p. 60) D-quantificational example, with an A-quantificational counterpart:

67. In Atlanta, there are many serial killers but not many unemployed men

68. In Atlanta, men are often serial killers but they are not often unemployed

67 is fine and involves free shifting with ‘many,’ but 68 is bizarre because ‘often’—being tied into the question under discussion—cannot freely shift in the way ‘many’ can.

Modals exhibit similar discourse-level constraints (Stanley 2005, p. 73; fn. 16). For instance:

69. Ann can speak Finnish but Ann can only speak English

sounds terrible, even though there is a sense in which Ann can speak Finnish (she has working vocal cords and the ability to learn a language) and a sense in which she cannot speak Finnish (she doesn’t know a word of it). This example involved a fixed “flavor” of dynamic modality concerning abilities. Matters seem even worse if one tries to shift flavors. For instance, it would sound utterly terrible for Ann to say:

70. Claire must be stealing the diamonds and must not be stealing the diamonds

⁵⁴ One invariantist who does run the good form of the semantic blindness objection in a way that would specifically challenge our semantics is Stanley (2005, pp. 52–55), in that Stanley denies semantic blindness for modals (at least in certain specific respects). We just disagree with Stanley about modals, siding with MacFarlane and others in thinking that all the aspects of semantic blindness are found with modals. (We think that once one has homophonic indirect reports, all the remaining aspects of semantic blindness come along for the ride.)

even though the first conjunct can be read as a true epistemic modal, and the second as a true deontic modal.⁵⁵

We remain neutral on whether adverbial quantifiers, modals, and ‘know’ exhibit the same degree of semantic blindness. Indeed we suspect that adverbial quantifiers exhibit less semantic blindness than modals, which in turn exhibit (slightly) less semantic blindness than ‘know.’ We leave open what might explain any such differences. Our point is only that some level of semantic blindness should be expected from A-quantifiers.

So we claim to have articulated a plausible contextualist semantics for ‘know,’ and thereby solved the semantic problem. Pending objections to our semantics, we can only conclude that, whatever problems *Epistemic contextualism* might face, lack of a plausible semantic implementation is not among them. (Even if our semantics should prove objectionable, we hope to have at least advanced the discussion by moving beyond the level of analogies, and displaying novel empirical considerations.)

5 Comparativism for other attitude verbs?

Assuming that our comparativist semantics is on the right track for ‘know,’ a natural question arises as to whether comparativism might be generalized to other attitude verbs. This question raises some interesting puzzles, interesting in part because they seem puzzling from any perspective (contextualist or not).

Once nice feature of the comparativist account of ‘always’ (Sect. 3.1) was how it fit a general template covering both adverbial quantifiers and modals. A second and related nice feature was how it explains a range of entailments. Holding fixed the domain restrictions, and just looking at universal and existential quantification, we find the following entailments:

must → *always* → *sometimes* → *can*

These entailments hold simply because the terms are making ever lesser comparative demands: from all possible situations, to all actual situations, to some actual situations, to some possible situation.

It would be lovely if our comparativist semantics as given in *Final pass* could likewise provide a more general template covering a range of other attitude verbs, in a way that can equally claim to explain a range of entailments among these verbs. Indeed—holding fixed domain restrictions—there seem to be at least the following entailments involving ‘know’⁵⁶:

⁵⁵ Because A-quantifiers take their contextual restrictors from the question under discussion of the discourse, and not from any syntactic material which might get bound when embedded, they also exhibit stability under modal and temporal embeddings. This difference between A-quantifiers and D-quantifiers also enables our semantics to avoid the embedding objections that Stanley (2005, pp. 111–113) offers to Lewis’s analogy with D-quantifiers.

⁵⁶ Williamson (2000, pp. 33–41) conjectures that knowledge is the most general factive stative attitude. If Williamson is right then *every* factive stative attitude entails knowledge.

The specific epistemics may be thought of as expressing ways of knowing, and the emotive factives may be thought of as expressing emotions from knowing. Thus if Ann sees that Claire stole the diamonds, or remembers that Claire stole the diamonds, or regrets that Claire stole the diamonds, or cares that Claire stole the diamonds, then she knows that Claire stole the diamonds; and if Ann knows that Claire stole the diamonds, then she believes that Claire stole the diamonds.

We have built in the entailment from knowledge to belief “by hand” in *Final pass*, simply by appending a ‘belief’-clause. But the entailments from the specific epistemics and emotive factives to knowledge still need explanation. It would be lovely if our comparativist semantics could provide a more general template covering these verbs, in a way that can explain the entailments to knowledge.

From our semantics for ‘know’ we can extract the general template:

$\llbracket ______ \rrbracket^c = \lambda E \lambda P \lambda x \lambda s.$ every s' such that s' is consistent with $______$, $C(s')$, and $E(s')$, is also such that $P(s')$; and $C(s)$; and x believes (P given C and E) on the basis of $______$

If we fill in the first blank with a specific epistemic verb, and the second and third blanks with anything at least as strong as ‘ x ’s evidence in s ,’ then we will get something that explains the knowledge entailment. For instance:

$\llbracket \text{remember} \rrbracket^c = \lambda E \lambda P \lambda x \lambda s.$ every s' such that s' is consistent with x ’s memorial evidence in s , $C(s')$, and $E(s')$, is also such that $P(s')$; and $C(s)$; and x believes (P given C and E) on the basis of x ’s memorial evidence in s

This fits the sense in which remembering is a way of knowing, by specifying the sort of evidence that eliminates the relevant alternatives—namely, memorial evidence. For ‘see,’ just replace each of the two occurrences of ‘memorial’ with ‘visual.’

For the emotive factives we might leave *Final pass* intact and append some sort of emotive clause (for which we will just use ‘ V -apt attitude’ as a placeholder):

$\llbracket \text{regret} \rrbracket^c = \lambda E \lambda P \lambda x \lambda s.$ every s' such that s' is consistent with x ’s evidence in s , $C(s')$, and $E(s')$, is also such that $P(s')$; and $C(s)$; and x believes (P given C and E) on the basis of x ’s evidence in s ; and x bears a regret-apt attitude in s to P

This fits the sense in which regretting involves an emotion from knowing, by treating regret as knowledge laden with a certain emotion. For other emotive factives, just adjust the final emotive clause accordingly.

So far this might suggest:

(*Comparativism generalized*) Specific epistemics and emotive factives all have an A-quantificational aspect, fitting our general template for ‘know’

Comparativism generalized would go some way towards fitting ‘know’ into a wider range of attitude verbs. It would explain the entailments listed above. And it would

also explain why the specific epistemics and emotive factives display hyperintensionality as well as the characteristic marks of A-quantification.

Indeed, question sensitivity and association with focus—our first two marks of A-quantification (Sects. 4.1, 4.2)—seem to line up perfectly with *Comparativism generalized*. The emotive factives are widely recognized as paradigm cases of association with focus (Beaver and Clark 2008, pp. 66–68). And the specific epistemics seem to have just as good a claim to question sensitivity and association with focus as does ‘know.’ Thus imagine that Ann and Ben, years afterwards, are reminiscing about Claire’s exploits.

(*WhoRem*) Claire stole the diamonds years ago. Ann and Ben are wondering who stole the diamonds back then, and Ann recalls finding Claire’s fingerprints on the safe. So Ann says:

71. I remember that Claire stole the diamonds

(*WhatRem*) Claire stole the diamonds years ago. Ann and Ben are wondering what Claire stole back then, and Ann recalls finding Claire’s fingerprints on the safe. So Ann says:

71. I remember that Claire stole the diamonds

Our intuitions are that the memorial ascription 71 is true in *WhoRem* but false in *WhatRem*, exactly as per the knowledge ascription 1 in *Who* versus *What*. Or:

(*WhoRemFocus*) Claire stole the diamonds years ago. Ann recalls finding Claire’s fingerprints on the safe. So Ann says:

72. I remember that CLAIRE stole the diamonds

(*WhatRemFocus*) Claire stole the diamonds years ago. Ann recalls finding Claire’s fingerprints on the safe. So Ann says:

73. I remember that Claire stole THE DIAMONDS

Our intuitions are that the memorial ascription 72 is true in *WhoRemFocus* but that 73 is false in *WhatRemFocus*, exactly as per the knowledge ascription 47 in *WhoFocus* versus 48 in *WhatFocus*.

Moreover, with other attitude verbs not covered by *Comparativism generalized* such as ‘believe,’ we seem not to find neither question sensitivity nor association with focus. Indeed, Beaver and Clark (2008, p. 51) argue for the generalization that an operator will show a merely pragmatic *quasi-association* with focus if it is a non-veridical propositional operator. So perhaps there is an even deeper generalization available, to a comparativist semantics for veridical propositional operators.

So far a lovely generalization seems to be emerging.⁵⁷ But when we consider scoping over restrictive clauses—our third mark of A-quantification (Sect. 4.3)—the

⁵⁷ Though there were already signs of trouble. In Sect. 2.2 we noted that ‘regret’ and other emotive factives are gradable but ‘know’ is not. There is also the fact that ‘regret’ and other emotive factives do

pattern collapses. In one direction, we find emotive factives that do not naturally take wide scope over ‘if’-clauses, such as ‘be thrilled.’ Thus imagine again that Ann and Ben are playing poker, and that Ann holds a pair of kings:

74. If Ben is holding a pair of queens Ann is thrilled she will win

74 is odd. It does not seem to have an interpretation on which Ann is being said to be thrilled that she will win in case Ben holds a pair of queens. Rather it only seems to have the interpretation on which, under the assumption that Ben is holding a pair of queens, Ann feels a thrill. Or imagine that Ann holds a pair of kings, and Ben has yet to receive his final cards:

75. If Ben will receive a pair of queens Ann is thrilled she will win

75 is bizarre, seeming to suggest that Ben’s receiving a pair of queens in the future could somehow act backwards in time and implant a present feeling of thrill in Ann. Or consider the following story:

(*ThrillCards*) Ann and Ben are playing poker, and Frank is watching. Frank circles behind Ann and sees that she holds a pair of kings. He thinks: “If Ben is holding a pair of queens Ann is thrilled she will win.” So he circles behind Ben and sees that Ben is indeed holding a pair of queens. So he thinks:

76. Ann is not yet thrilled that she will win, but she will win

Frank’s thoughts—if one can look past the oddity of the first of his thoughts—seem blatantly incoherent in *ThrillCards*. Putting this together, it seems that ‘be thrilled’ cannot take wide scope over an ‘if’-clause.

In the other direction, we find that ‘believe’ does seem to interact with restrictive clauses in the way ‘know’ does, as seen in:

77. If Ben holds a pair of queens Ann believes she will win

77 seems to us to be most naturally read with ‘believe’ taking wide scope, surface form notwithstanding. Or imagine that Ann holds a pair of kings, and Ben has yet to receive his final cards. Now consider:

78. If Ben will receive a pair of queens Ann believes she will win

78 is naturally read with ‘believe’ taking wide scope. Likewise consider:

(*BelieveCards*) Ann and Ben are playing poker, and Frank is watching. Frank circles behind Ann and sees that she holds a pair of kings. He thinks: “If Ben is holding a pair of queens Ann believes she will win.” So he circles behind Ben and sees that Ben is indeed holding a pair of queens. So he thinks:

Footnote 57 continued

not take *whether*-complements but ‘know’ does. This generalization looks to hold cross-linguistically, suggesting an underlying semantic difference.

79. Ann does not yet believe that she will win, but she will win

Frank's thoughts seem coherent in *BelieveCards*, which requires the first of his thoughts to be read with 'believe' taking wide scope.

Accordingly we must remain agnostic as to *Comparativism generalized*, and leave the prospect of generalizing comparativism to other attitude verbs unsettled. The attitude verbs seem to exhibit complex patterns resisting simple generalization. These patterns should be a puzzle for everyone, contextualist or not. The puzzle is how (i) question sensitivity and association with focus could disassociate from (ii) scoping over restrictive clauses (as seems to happen in both directions, with 'be thrilled' having the former but not the latter features, and 'believe' having the latter but not the former features).

Perhaps we are seeing distinct types of A-quantificational structures? Suppose that question sensitivity and association with focus arise from the presence of a contextual restrictor argument (anaphoric on the question under discussion), and scoping over restrictive clauses arises from the presence of an explicit restrictor argument. Then perhaps some A-quantificational structures—namely those like 'be thrilled' that display question sensitivity and association with focus but do not tend to scope over restrictive clauses—have contextual restrictor arguments but lack explicit restrictor arguments. And perhaps other A-quantificational structures—namely those like 'believe' that do tend to scope over restrictive clauses but do not display question sensitivity or association with focus—have explicit restrictor arguments but lack contextual restrictor arguments. And perhaps still other A-quantificational structures—namely those like 'always' and 'know' that display question sensitivity and association with focus as well as a tendency to scope over restrictive clauses—do so because they have both contextual restrictor arguments as well as explicit restrictor arguments. We leave the question open.

Acknowledgments Thanks to George Bealer, David Beaver, Jessica Brown, Stewart Cohen, Keith DeRose, Kai von Fintel, Georgi Gardiner, Michael Glanzberg, Brendan Balcerak Jackson, Dirk Kindermann, Joshua Knobe, Justin Khoo, Angelika Kratzer, Nick Kroll, Carlotta Pavese, Craige Roberts, Jason Stanley, Tamina Stephenson, Bruno Whittle, and audiences at the ANU Epistemology Workshop, the Arché Research Centre, NYU Linguistics, the Paris Contextualism and Compositionality Conference, the Yale Friday faculty lunch, the 6th International Symposium of Logic, Language, and Communication in Rigá, and the Cornell Workshop in Linguistics & Philosophy.

References

- Aloni, M., & Égré, P. (2008). Alternative questions and knowledge attributions. *Philosophical Quarterly*, 60, 1–27.
- Anand, P., & Nevins, A. (2004). Shifty operators in changing contexts. In K. Watanabe & R. B. Young (Eds.), *Proceedings of SALT XIV* (pp. 20–37). Ithaca, NY: CLC Publications.
- AnderBois, S. (2009). *If*-suspenders and the existential presupposition of questions. Accessed March 31, 2013, from <https://sites.google.com/site/anderbois2/Research/CUSP2Handout.pdf>.
- Austin, J. L. (1946). Other minds. *Proceedings of the Aristotelian Society*, 20, 149–187.
- Bach, K. (1994). Conversational implicature. *Mind and Language*, 9, 124–162.
- Bach, K. (2005). The Emperor's new 'knows'. In G. Preyer & G. Peter (Eds.), *Contextualism in philosophy: Knowledge, meaning, and truth* (pp. 51–89). New York: Oxford University Press.

- Beaver, D., & Clark, B. (2008). *Sense and sensitivity: How focus determines meaning*. Oxford: Wiley-Blackwell.
- Berman, S. (1987). Situation-based semantics for adverbs of quantification. In J. Blevins & A. Vainakka (Eds.), *University of Massachusetts occasional papers in linguistics 12: Issues in semantics* (pp. 45–68). Amherst, MA: GLSA Publications.
- Blome-Tillmann, M. (2008). The indexicality of ‘knowledge’. *Philosophical Studies*, 138, 29–53.
- Blome-Tillmann, M. (2009). Knowledge and presuppositions. *Mind*, 118, 241–294.
- Buckwalter, W. (2010). Knowledge isn’t closed on Saturday: A study in ordinary language. *Review of Philosophy and Psychology*, 1, 395–406.
- Cappelen, H., & LePore, E. (2005). *Insensitive semantics: A defense of semantic minimalism and speech act pluralism*. Oxford: Blackwell.
- Carlson, L. (1983). *Dialogue games: An approach to discourse analysis*. Dordrecht: Reidel.
- Carston, R. (2002). *Thoughts and utterances: The pragmatics of explicit communication*. Oxford: Blackwell.
- Cohen, S. (1986). Knowledge and context. *Journal of Philosophy*, 83, 574–583.
- 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.
- Cohen, S. (1999). Contextualism, skepticism, and the structure of reasons. *Philosophical Perspectives*, 13, 57–89.
- Collingwood, R. G. (1924). *Speculum mentis*. Oxford: Oxford University Press.
- DeRose, K. (1992). Contextualism and knowledge attributions. *Philosophy and Phenomenological Research*, 52, 913–929.
- DeRose, K. (1995). Solving the skeptical problem. *Philosophical Review*, 104, 1–52.
- DeRose, K. (2008). Gradable adjectives: A defence of pluralism. *Australasian Journal of Philosophy*, 86, 141–160.
- DeRose, K. (2009). *The case for contextualism: Knowledge, skepticism, and context* (Vol. 1). New York: Oxford University Press.
- Dretske, F. (1970). Epistemic operators. *Journal of Philosophy*, 67, 1007–1023.
- Dretske, F. (1981). The pragmatic dimensions of knowledge. *Philosophical Studies*, 40, 363–378.
- Dretske, F. (1991). Knowledge: Sanford and Cohen. In B. McLaughlin (Ed.), *Dretske and his critics* (pp. 185–196). New York: Oxford University Press.
- Dutant, J. (2007). The limited gradability of “knows”. Accessed August 8, 2010, from http://julien.dutant.free.fr/papers/JDutant_TheLimitedGradabilityOfKnow.pdf.
- Egan, A., Hawthorne, J., & Weatherston, B. (2005). Epistemic modals in context. In G. Preyer & G. Peter (Eds.), *Contextualism in philosophy: Knowledge, meaning, and truth* (pp. 131–168). Oxford: Oxford University Press.
- Elbourne, P. (2005). *Situations and individuals*. Cambridge, MA: MIT Press.
- Fantl, J., & McGrath, M. (2002). Evidence, pragmatics, and justification. *Philosophical Review*, 111, 67–94.
- Feltz, A., & Zarpentine, C. (2010). Do you know more when it matters less? *Philosophical Psychology*, 23, 683–706.
- Frank, A. (1996). *Context dependence and modal constructions*. Ph.D. Dissertation, University of Stuttgart.
- Gillies, A. (2010). Iffiness. *Semantics and Pragmatics*, 3, 1–42.
- Ginzburg, J. (1996). The semantics of interrogatives. In S. Lappin (Ed.), *The handbook of contemporary semantic theory* (pp. 385–422). Oxford: Basil Blackwell.
- Goldman, A. (1976). Discrimination and perceptual knowledge. *Journal of Philosophy*, 73, 771–791.
- Groenendijk, J., & Roelofsen, F. (2009). Inquisitive semantics and pragmatics. In *Proceedings of SPR-09, ILCLI international workshop on semantics, pragmatics, and rhetoric* (pp. 41–72). San Sebastián: University of the Basque Country Press.
- Groenendijk, J., & Stokhof, M. (1984). Questions. In J. van Benthem & A. ter Meulen (Eds.), *Handbook of logic and language* (pp. 1055–1124). Amsterdam: Elsevier.
- Hamblin, C. L. (1973). Questions in Montague English. *Foundations of Language*, 10, 41–53.
- Hawthorne, J. (2004). *Knowledge and lotteries*. Oxford: Oxford University Press.
- Heim, I. (1990). E-type pronouns and donkey anaphora. *Linguistics and Philosophy*, 13, 137–177.
- Heim, I. (1992). Presupposition projection and the semantics of attitude verbs. *Journal of Semantics*, 9, 183–221.

- Heller, M. (1999). The proper role for contextualism in an anti-luck epistemology. *Philosophical Perspectives*, 13, 115–129.
- Herburger, I. (2000). *What counts: Focus and quantification*. Cambridge, MA: MIT Press.
- Hintikka, J. (1962). *Knowledge and belief: An introduction to the logic of the two notions*. Ithaca, NY: Cornell University Press.
- 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.
- Horn, L. (1972). *On the semantic properties of logical operators in English*. Ph.D. Dissertation, UCLA.
- Ichikawa, J. (2011). Quantifiers and epistemic contextualism. *Philosophical Studies*, 155, 383–398.
- Kaplan, D. (1989). Demonstratives. In J. Perry, J. Almog, & H. Wettstein (Eds.), *Themes from Kaplan* (pp. 481–563). Oxford: Oxford University Press.
- Kennedy, C. (2007). Vagueness and grammar: The semantics of relative and absolute gradable adjectives. *Linguistics and Philosophy*, 30, 1–45.
- Kennedy, C., & McNally, L. (2005). Scale structure and the semantic typology of gradable predicates. *Language*, 81, 345–381.
- Khoo, J. (2011). Operators or restrictors? A reply to Gillies. *Semantics and Pragmatics*, 4, 1–25.
- Kiss, K. (1998). Identificational focus versus information focus. *Language*, 74, 245–273.
- Kratzer, A. (1989). An investigation of the lumps of thought. *Linguistics and Philosophy*, 12, 607–653.
- Kratzer, A. (1991). Modality. In A. von Stechow & D. Wunderlich (Eds.), *Semantics: An international handbook of contemporary research* (pp. 639–650). Berlin: Walter de Gruyter.
- Kratzer, A. (2009). Situations in natural language semantics. *Stanford Encyclopedia of Philosophy*. Accessed August 8, 2010, from <http://plato.stanford.edu/entries/situations-semantics/>.
- Kratzer, A. (2012). *Collected papers on modals and conditionals*. Oxford: Oxford University Press.
- Lewis, D. (1975). Adverbs of quantification. In E. Keenan (Ed.), *Formal semantics of natural language*. Cambridge: Cambridge University Press.
- Lewis, D. (1979). Scorekeeping in a language game. *Journal of Philosophical Logic*, 8, 339–359.
- Lewis, D. (1986). *On the plurality of worlds*. Oxford: Blackwell.
- Lewis, D. (1996). Elusive knowledge. *Australasian Journal of Philosophy*, 74, 549–567.
- Lihoreau, F. (2008). Relevant alternatives contextualism and ordinary contingent knowledge. *Disputatio*, 24, 281–294.
- Ludlow, P. (2008). Cheap contextualism. *Philosophical Issues*, 18, 104–129.
- MacFarlane, J. (2005). The assessment-sensitivity of knowledge attributions. *Oxford Studies in Epistemology*, 1, 235–271.
- MacFarlane, J. (2009). Nonindexical contextualism. *Synthese*, 166, 231–250.
- MacFarlane, J. (2011). Epistemic modals are assessment-sensitive. In A. Egan & B. Weatherson (Eds.), *Epistemic modality* (pp. 144–178). Oxford: Oxford University Press.
- Matthewson, L., Rummelmann, H., & Davis, H. (2005). Modality in St'at'imcets. In *Papers for the 40th international conference on Salish and neighboring languages* (pp. 166–183). Vancouver: University of British Columbia Linguistics.
- May, J., Sinnott-Armstrong, W., Hull, J. G., & Zimmerman, A. (2010). Practical interests, relevant alternatives, and knowledge attributions: An empirical study. *Review of Philosophy and Psychology*, 1, 265–273.
- Neta, R. (2002). S Knows that P. *Noûs*, 36, 663–681.
- Neta, R. (2003). Skepticism, contextualism, and semantic self-knowledge. *Philosophy and Phenomenological Research*, 67, 396–411.
- Partee, B. (1995). Quantificational structures and compositionality. In E. Bach, E. Jelinek, A. Kratzer, & B. Partee (Eds.), *Quantification in natural languages* (pp. 541–601). Dordrecht: Kluwer.
- Pavese, C. (manuscript). Degrees of knowledge how.
- Percus, O. (2000). Constraints on some other variables in syntax. *Natural Language Semantics*, 8, 173–229.
- Potts, C. (2011). On the negativity of negation. In D. Lutz & N. Li (Eds.), *Proceedings of SALT 20* (pp. 636–659). Ithaca, NY: CLC Publications.
- Quine, W. V. O. (1966). *Elementary logic* (Revised ed.). Cambridge, MA: Harvard University Press.
- Roberts, C. (1996). Information structure in discourse: Towards an integrated formal theory of pragmatics. In J.-H. Yoon & A. Kathol (Eds.), *Working papers in linguistics no. 49: Papers in semantics* (pp. 91–136). Columbus, OH: Ohio State University.
- Roberts, C. (2004). Context in dynamic interpretation. In L. R. Horn & G. Ward (Eds.), *The handbook of pragmatics* (pp. 197–220). Oxford: Blackwell.

- Rooth, M. (1996). Focus. In S. Lappin (Ed.), *The handbook of contemporary semantic theory* (pp. 271–298). Oxford: Blackwell.
- Rysiew, P. (2006). Motivating the relevant alternatives approach. *Canadian Journal of Philosophy*, 36, 259–279.
- Sanford, D. (1991). Proper knowledge. In B. McLaughlin (Ed.), *Dretske and his critics* (pp. 38–51). Oxford: Oxford University Press.
- Schaffer, J. (2004). From contextualism to contrastivism. *Philosophical Studies*, 119, 73–103.
- Schaffer, J. (2005a). Contrastive knowledge. *Oxford Studies in Epistemology*, 1, 235–271.
- Schaffer, J. (2005b). What shifts? Thresholds, standards, or alternatives? In G. Preyer & G. Peter (Eds.), *Contextualism in philosophy* (pp. 115–130). Oxford: Oxford University Press.
- Schaffer, J. (2007). Knowing the answer. *Philosophy and Phenomenological Research*, 75, 383–403.
- Schaffer, J. (2011). Perspective in taste predicates and epistemic modals. In A. Egan & B. Weatherson (Eds.), *Epistemic modality* (pp. 179–226). Oxford: Oxford University Press.
- Schaffer, J., & Knobe, J. (2011). Contrastive knowledge surveyed. *Noûs*, 46(4), 675–708.
- Schiffer, S. (1996). Contextualist solutions to scepticism. *Proceedings of the Aristotelian Society*, 96, 317–333.
- Schlenker, P. (2003). A plea for monsters. *Linguistics and Philosophy*, 26, 29–120.
- Schoubye, A. (2009). Descriptions, truth value intuitions, and questions. *Linguistics and Philosophy*, 32, 583–617.
- Simons, M., Tonhauser, J., Beaver, D., & Roberts, C. (2011). What projects and why. In D. Lutz & N. Li (Eds.), *Proceedings of SALT 20* (pp. 309–327). Ithaca, NY: CLC Publications.
- Soames, S. (2008). Drawing the line between meaning and implicature—And relating both to assertion. *Noûs*, 42, 440–465.
- Sperber, D., & Wilson, D. (1986). *Relevance: Communication and cognition*. Cambridge, MA: Harvard University Press.
- Stainton, R. (2010). Contextualism in epistemology and the context sensitivity of ‘knows’. In J. K. Campbell, M. O’Rourke, & H. Silverstein (Eds.), *Knowledge and skepticism: Topics in contemporary philosophy* (Vol. 5, pp. 113–139). Cambridge, MA: MIT Press.
- Stanley, J. (2005). *Knowledge and practical interests*. Oxford: Oxford University Press.
- Stephenson, T. (2007). Judge dependence, epistemic modals, and predicates of personal taste. *Linguistics and Philosophy*, 30, 487–525.
- Stine, G. C. (1976). Scepticism, relevant alternatives, and deductive closure. *Philosophical Studies*, 29, 249–261.
- Szabó, Z. G. (2000). Descriptions and uniqueness. *Philosophical Studies*, 101, 29–57.
- Szabó, Z. G. (2005a). The loss of uniqueness. *Mind*, 114, 1185–1222.
- Szabó, Z. G. (2005b). Introduction. In Z. G. Szabó (Ed.), *Semantics versus pragmatics* (pp. 1–14). Oxford: Oxford University Press.
- Szabó, Z. G. (2006). Sensitivity training. *Mind and Language*, 21, 31–38.
- Szabó, Z. G. (2011). Bare quantifiers. *Philosophical Review*, 120, 247–283.
- Unger, P. (1975). *Ignorance: A case for scepticism*. Oxford: Oxford University Press.
- von Fintel, K. (1994). *Restrictions on quantifier domains*. Ph.D. Dissertation, University of Massachusetts-Amherst.
- von Fintel, K. (2004). A minimal theory of adverbial quantification. In H. Kamp & B. Partee (Eds.), *Context-dependence in the analysis of linguistic meaning* (pp. 137–175). Amsterdam: Elsevier.
- von Stechow, A. (1982). Structured propositions. Technical Report 59, Sonderforschungsbereich 99, University of Konstanz.
- von Stechow, A. (2004). Binding by verbs: Tense, person, and mood under attitudes. In H. Lohnstein & S. Trissler (Eds.), *The syntax and semantics of the left periphery* (pp. 431–488). Berlin: Mouton de Gruyter.
- Weatherson, B. (2005). True, truer, truest. *Philosophical Studies*, 123, 47–70.
- Williamson, T. (2000). *Knowledge and its limits*. Oxford: Oxford University Press.
- Williamson, T. (2005). Contextualism, subject-sensitive invariantism and knowledge of knowledge. *Philosophical Quarterly*, 55, 213–235.
- Wittgenstein, L. (1969). *On certainty* (G. E. M. Anscombe & G. H. von Wright, Eds.). New York: Harper & Row.
- Yalcin, S. (2007). Epistemic modals. *Mind*, 116, 983–1026.
- Zvolenszky, Z. (2002). Is a possible worlds semantics for modality possible? In B. Jackson (Ed.), *Proceedings of SALT 12* (pp. 339–359). Ithaca, NY: CLC Publications.