

OXFORD

Contextualism in Philosophy

Knowledge, Meaning, and Truth

edited by

GERHARD PREYER
AND GEORG PETER

What Shifts? Thresholds, Standards, or Alternatives?

JONATHAN SCHAFFER

I. What Shifts?

Much of the extant discussion of epistemic contextualism focuses on the question of whether contextualism resolves skeptical paradoxes.¹ Understandably. Yet there has been less discussion as to the internal structure of contextualist theories. Regrettably. Here, for instance, are two questions that could stand further discussion: (i) what is the *linguistic basis* for contextualism, and (ii) what is the *parameter that shifts* with context?

The question of linguistic basis can be understood as a request for an explanation of how the truth conditions for knowledge ascriptions are supposed to shift with context. Is there some extra *variable* hidden in the syntax? Is 'knows' to be treated as a semantic *indexical*? Are there general *indices* of semantic evaluation that impact knowledge ascriptions? Or...? (Here the range of options will depend on one's overall view of the levels of linguistic structure—further discussion of this issue is under way in this volume.)

The question of what parameter shifts can be understood as a request for an explanation of which epistemic gear the wheels of context turn. Is there some shifting *threshold* for justification? Is there a shifting *standard* of epistemic

¹ Including such foundational presentations of contextualism as Stewart Cohen 1988 and 1999, Keith DeRose 1995, David Lewis 1996, and Mark Heller 1999.

position? Or is there a shifting set of epistemic *alternatives*? Or...? (Here it is not even *prima facie* clear how or whether these options differ, though I will clarify the differences below.)

In what follows, I will focus on the question of what parameter shifts with context. In section 2, I will display four desiderata for an answer to what shifts. In sections 3–5, I will consider thresholds, standards, and alternatives (respectively) in light of these desiderata, and uphold alternatives as the parameter of shift. In section 6, I will cast a parting glance at the linguistic basis for contextualism.

(While the discussion to follow will presuppose that some form of contextualism is true, it may still be of interest to invariantists who want to shunt contextual dependency into the pragmatics. For invariantists also face the question of what shifts with context, in developing their pragmatic account.)

2. Desiderata

What parameter shifts with context? Before considering candidate answers to this question, it will prove useful to display some desiderata for the candidates to meet. I offer four (interrelated) desiderata.

First, the parameter of shift must be *linguistically plausible*. That is, the alleged parameter associated with ‘knows’ should be a linguistically general parameter, associated with a natural class of expressions of which ‘knows’ is an instance. In other words, it will not do to invent a special parameter just for ‘knows’, or to import one from an unrelated class of expressions. Thus:

D1. What shifts should be a linguistically general parameter, associated with a natural class of expressions of which ‘knows’ is an instance.

Second, the parameter of shift must be *predictively adequate*. That is, the alleged parameter associated with ‘knows’ should shift in ways that match intuitions about the acceptability of knowledge ascriptions. In other words, it will not do to associate a parameter whose shifts are triggered by considerations such as whether it is Wednesday, or any considerations different from those that trigger shifts in intuitions about knowledge ascriptions. Thus:

D2. What shifts should sway with intuitions about the acceptability of knowledge ascriptions.

Third, an answer to what shifts must be *skeptically resolving*. That is, the alleged parameter associated with ‘knows’ should shift in ways that vindicate

contextualist solutions to both moderate and radical skepticism. To do so, it should render most ordinary knowledge ascriptions true in ordinary contexts, some (those associated with the specific doubts in play) false in moderately skeptical contexts, and most (or perhaps all) false in radically skeptical contexts. For instance, in ordinary contexts the following sorts of claims should count as true (the reader should fill in the background details in the obvious ways): (i) ‘I know that my car is parked on Elm’, and (ii) ‘I know that the movie starts at nine’. In a moderately skeptical context in which unresolved doubts have been raised as to whether my car has been stolen and relocated, (i) should count as false, though (ii) should still count as true (no doubts have yet been raised about *that*). Whereas in a radically skeptical context in which unresolved doubts have been raised as to whether one is dreaming, or a brain-in-a-vat, etc., (i) and (ii) should both count as false. Thus:

D3. What shifts should vindicate contextualist treatments of both moderate and radical skepticism.

Fourth, an answer to what shifts must *illuminate inquiry*. That is, the alleged parameter associated with ‘knows’ should connect to the practical role that knowledge ascriptions play within the larger project of inquiry. The practical role of knowledge ascriptions is (at least in part) to certify that the subject can answer the question.² To connect to this role, the parameter must be capable of scoring the question. For instance, consider the following inquiries: (i) ‘Is there a goldfinch in the garden, or a blue jay?’, (ii) ‘Is there a goldfinch in the garden, or a canary?’, and (iii) ‘Is there a goldfinch in the garden, or at the neighbor’s?’ The role of an ascription of ‘I know that there is a goldfinch in the garden’ differs in the context of (i)–(iii). With (i), such an ascription certifies that the speaker can tell a goldfinch from a blue jay; with (ii), it certifies that the speaker can tell a goldfinch from a canary (a harder task); while with (iii), it certifies

² Thus Christopher Hookway remarks: “The central focus of epistemic evaluation is . . . the activity of inquiry . . . When we conduct an inquiry . . . we attempt to formulate questions and to answer them correctly.” (1996, p. 7) And Hector-Neri Castañeda maintains: “[K]nowledge involves essentially the non-doxastic component of a power to answer a question.” (1980, p. 194) Connections between knowledge ascriptions and the ability to answer questions emerge in our practice of testing students and fielding questions. Thus the professor may preface the test with: “Let’s see what you know”. One may field a question with “I know”, or pass it with “Ask Pam, she knows”. Such connections emerge most directly with knowledge-*wh* ascriptions. Thus James Higginbotham suggests that the sentence: “Mary knows who John saw” should be interpreted as: “Mary knows the (or an) answer to the question who John saw.” (1993, p. 205) For further discussion, see Schaffer *manuscript*.

that the speaker can to tell the garden from the neighbor's (an entirely different task). The parameter of shift should explain how this can happen. Thus:

D4. What shifts should illuminate the role of knowledge ascriptions in our practices of inquiry, by keeping score of the question.

I think (D1)–(D3) should be relatively uncontroversial among contextualists. Perhaps (D4) will be somewhat controversial, if only because the role of knowledge ascriptions in inquiry is not so well explored. But never mind. For I will argue that thresholds and standards parameters fail all of (D1)–(D4), while an alternatives parameter satisfies them all.

3. Thresholds

So what parameter shifts with context? Let me begin by considering one possible answer (suggested by some remarks in Cohen, 1988), according to which what shifts is the *threshold* required for 'justified'. More precisely, for a subject s with a belief p , s is assigned an absolute degree d of justification for p . What shifts is whether d suffices for 'justified'. Thus:

T. What shifts is the threshold of justification sufficient for 'justified'.

Picturesquely, think of degrees of justification as measured on the interval $[0, 1]$, and think of context as selecting a threshold t for 'justified' on this interval. What shifts is whether $d < t$ or $d \geq t$. See Figure 5.1. Here s 's belief that p counts as 'justified' in context₁, 'unjustified' in context₂.




Figure 5.1

(T) derives its plausibility from the following observations: (i) justification is a necessary condition for knowledge, (ii) justification is a *vague* notion (or at least, 'justified' bears all the hallmarks of a vague predicate), and (iii) vague notions generally have contextually variable thresholds. Thus, a given subject s has an absolute degree h of height, an absolute degree w of income, etc. What shifts is whether s 's degree of height h suffices for 'tall', whether s 's degree of income w suffices for 'rich', etc.³ So far, so good.

³ This is an oversimplification. For a detailed account see Christopher Kennedy 1999.

But does (T) respect desiderata (D1)–(D4)? I am afraid that the answer is *no* in each case. Starting with (D1), while a linguistically general parameter has been identified, it is *not* a parameter associated with expressions of which ‘knows’ is an instance. It is rather a parameter associated with *gradable adjectives* like ‘tall’, ‘rich’, and ‘justified’. And ‘knows’ is not an adjective, much less a gradable one (Stanley, *forthcoming*).

The fact that a threshold parameter is associated with ‘justified’ does not help, as that *predicate* does not occur in knowledge ascriptions. Thus the vagueness it bears should not be triggered. Knowledge may well entail justification, but that does not render ‘knows’ itself vague. After all, possessing precisely zero hairs entails baldness, but that hardly renders ‘possesses precisely zero hairs’ vague. The mere conceptual entailments of a term play no role here.

Turning to (D2), (T) predicts the wrong shifts. That is, (T) predicts that we would flit from ‘knows’ to ‘does not know’ and back, as the bar of justification rises and falls. Think of how we may flit from ‘rich’ to ‘not rich’ and back. Here we may invoke *comparison classes* (‘he, like all professors, is not rich’ and ‘he, like all Americans, is rich’), or we may simply *draw the line* (‘I mean: at least a millionaire’). But none of this comparing or line drawing seems to trigger any shifts with ‘knows’. Instead, what seems to get us to shift from ‘knows’ to ‘does not know’ is the invocation of *specific doubts*. And what seems to get us to shift back to ‘knows’ is forgetting such doubts entirely.⁴ The pattern of shift for ‘knows’ does not match the overall pattern of a shifting threshold.

Moving to (D3), (T) does not fit contextualist solutions to skepticism, for two reasons. First, it is unclear why raising skeptical doubts should generate *any shift at all* in the threshold. Why should raising doubts about whether one is a brain-in-a-vat have any impact on the threshold at all, much less drive it to the max? By analogy, it would be as if the mention of a humanly unreachable height (say, 1 mile) would drive the threshold for ‘tall’ through the roof. *That* would be surprising.

Second, and most crucially, raising doubts would shift the threshold in the wrong way for *moderate skepticism*. When thresholds shift, they do so in ways that *globally infect* other truth-values in that context. For instance, raising the bar of tallness for anyone raises it for everyone. If *x* does not satisfy ‘tall’ in

⁴ As David Lewis 1979 points out, there is an asymmetry between (i) the *ease* by which the skeptic can shift to “does not know”, and (ii) the *difficulty* for the dogmatist in shifting back to “knows”. Skeptical doubts do not dissipate until the conversation is forgotten.

context c , and y has a height less than or equal to x 's, then y cannot satisfy 'tall' in c either.⁵

Yet such global infection does not occur in moderately skeptical scenarios. Here our doubts are *localized*. For instance, if unresolved doubts are raised as to whether one's car has been stolen and relocated, car-location knowledge claims should now count as false, though movie-schedule knowledge claims should still count as true (§2). Conversely, if unresolved doubts are raised as to whether the movie schedule has been misprinted, then car-location knowledge claims should still count as true, though movie-schedule knowledge claims should now count as false. But with thresholds, this localization of doubt cannot happen. These claims will be assigned absolute degrees d_1 and d_2 of justification. If $d_1 = d_2$ (as is roughly plausible), then the threshold cannot be raised past d_1 without passing d_2 as well. Whereas if $d_1 \neq d_2$, then the threshold cannot be raised past whichever is greater without passing the other as well. In general, (T) renders local skeptical scenarios overly global. It means that raising the bar of justification anywhere raises it everywhere.⁶

Shifting finally to (D₄), (T) does not keep score of the question, for two reasons. First, it is unclear that there is any relation at all between a line of inquiry and a threshold of justification. Why should the threshold of justification for 'I know that there is a goldfinch in the garden' respond at all to whether the inquiry is between (i) a goldfinch and a blue jay, (ii) a goldfinch and a canary, or (iii) the garden and the neighbor's? There seems to be no connection here.

Second and most crucially, the globality of thresholds would conflate distinct lines of inquiry. For instance, in inquiry (i) it is presupposed that the bird is not a canary, and that it is in the garden; in (ii) it is presupposed that the bird is not a blue jay, and (still) that it is in the garden; while in (iii) it is presupposed that the bird is a goldfinch (*a fortiori* neither a blue jay nor a canary), but no longer presupposed that it is in the garden. But thresholds make no such distinctions. 'Knows' will require some threshold t_1 in (i), t_2 in (ii), and t_3 in (iii). Whichever one or more of the thresholds is highest will simply subsume the others. The use of a thresholds parameter will conflate resolving the inquiry/inquiries associated with the highest threshold(s) with

⁵ Infection is related to the phenomenon of *penumbral connection*, in which judgments about borderline cases are related. See Kit Fine 1975 for further discussion.

⁶ Nor can the contextualist hope that we can shift contexts rapidly enough to get the effects of localized doubts. For it is crucial to the contextualist explanation of why dogmatic claims ring false in skeptical contexts (rather than forcing accommodation so they ring true), that skeptical doubts do not dissipate easily (see footnote 5).

resolving all the other inquiries associated with equal or lesser thresholds, regardless of any differences in presupposition. (T) thus submerges specific differences in what is posed and what is presupposed, under a general bar of justification.

There is an overall moral to be drawn. What shifts with ‘knows’ needs to be *locally responsive to specific doubts and questions*. Thresholds are too monolithic.

4. Standards

I now turn to a second possible answer (developed by DeRose, 1995, and endorsed by Heller, 1999), according to which what shifts is the *strength of epistemic position* required for knowledge. More precisely, for a subject s with a belief p , s is assigned an absolute strength of position r for p relative to similarity metric m , where r is the maximal radius in logical space as ordered by m through which s can *track* the truth or falsity of p .⁷ What shifts is both which similarity metric is in play, and whether r extends far enough for ‘knows’ on that metric. Thus:

S. What shifts is the metric of similarity, together with the standard of how far one must track for ‘knows’.

Picturesquely, think of logical space as ordered into nested spheres via the similarity metric m , and think of context as both selecting m and selecting a standards radius l out to which one must track to satisfy ‘knows’. What shifts is both the metric m , and whether $r < l$ or $r \geq l$ given m . Thus consider a toy model in which s can track the truth at α and w_1 , but not at w_2 . Suppose that context₁ sets m to $\langle \alpha, w_1, w_2 \rangle$ and sets l to 1 (see Figure 5.2). Then s ’s belief that p counts as ‘knowledge’ in context₁. But suppose that context₂, while keeping m at $\langle \alpha, w_1, w_2 \rangle$, sets l to 2 (see Figure 5.3). Then s ’s belief that p does not count as ‘knowledge’ in context₂. Or suppose that context₃ sets m to $\langle \alpha, w_2, w_1 \rangle$, while keeping l at 1 (see Figure 5.4). Then s ’s belief that p does not count as ‘knowledge’ in context₃.

Standards may seem thematically similar to thresholds, if one assumes that s ’s position r (how far s can track p) is s ’s degree of justification d . But one

⁷ The notion of tracking is borrowed from Robert Nozick 1981. The idea is that s tracks p through a sphere of radius r iff, for all worlds within r , s is right about p (that is, if p is true, then s believes it; if p is false, then s disbelieves it).

need not assume this, and should not, at least if one has anything like an *internalist* conception of justification.⁸ One should not assume that distance in logical space and degree of justification will generate anything like the same orderings.

In any case, (T) and (S) are formally distinct in the following two ways. First, the degree of justification has an *upper bound* of 1, perhaps enjoyed in cases like the mathematician's belief that $2+2=4$. Whereas the position r presumably has no upper bound. Second, (S) involves an *additional parameter* of contextual variation, namely the similarity metric m .

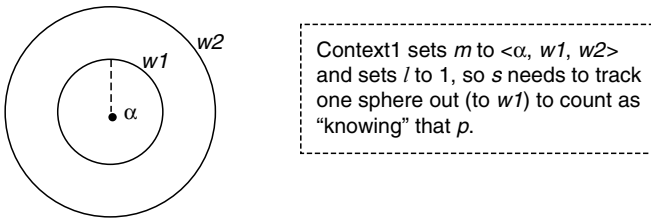


Figure 5.2

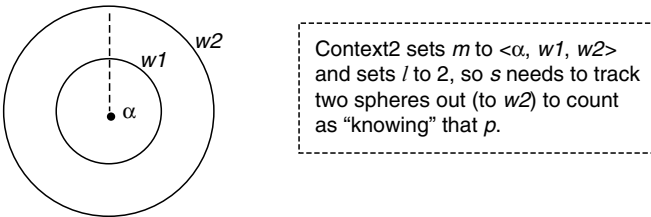


Figure 5.3

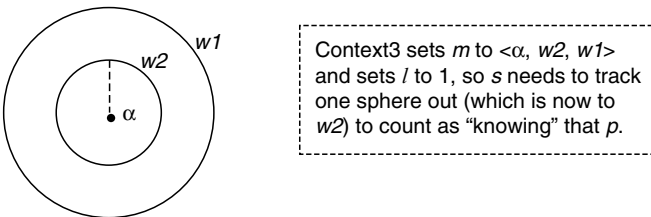


Figure 5.4

⁸ By an internalist conception of justification, I mean one according to which a justification must be in principle accessible to s 's conscious reflection. For further discussion see Richard Feldman and Earl Conee 1985, and Roderick Chisholm 1988, *inter alia*.

(S) derives its plausibility from the following observations: (i) knowledge entails an ability to track the truth,⁹ (ii) how far out into logical space one tracks the truth may vary, such that (iii) there might be a parameter in the language corresponding to such variation. So far, so good. But does (S) respect desiderata (D₁)–(D₄)? I am afraid the answer is *no* in each case. Starting with (D₁), no linguistically *general* parameter has been identified, much less one associated with expressions of which ‘knows’ is an instance. There seems to be no *precedent* for this form of parameter in the language. It seems a pure invention.

Turning to (D₂), (S) predicts the wrong shifts. That is, (S) predicts that we would flit from ‘knows’ to ‘does not know’ and back, as the metric *m* contorts and as the standard *l* expands and contracts. But think of the sort of maneuvers that are commonly taken to contort the similarity metric, namely the emphasizing of certain *respects of similarity* as salient. That is, if one emphasizes the presence of the atom bomb, one will naturally take the closest world at which Caesar is a general in Korea as a world *w₁* in which Caesar uses the bomb. While if one emphasizes the beliefs of Romans, one will naturally take the closest world at which Caesar is a general in Korea as a world *w₂* in which Caesar uses catapults. But none of this weighing of respects seems to accomplish so much with ‘knows’. Instead, what seems to get us to shift between ‘knows’ and ‘does not know’ is the invocation and revocation of specific doubts.¹⁰

Likewise think of the sorts of maneuvers that might expand and contract the standard of epistemic strength. Presumably, if anything can impact the standard, assertions of ‘that is enough’ and ‘that is not enough’ should do (though here the lack of a linguistic precedent makes it problematic to guess at which maneuvers would be appropriate). But once again such line drawing does not seem to accomplish much with ‘knows’. Again, what does the work are specific doubts.

⁹ Of course the tracking view of knowledge is highly contentious. Advocates of standards such as DeRose speak of tracking as being “at least roughly correct.” (1995, p. 25) For present purposes I will not contest this. I should also add that (S) may help add to the plausibility of the tracking view, by reconciling it with closure principles (DeRose 1995).

¹⁰ How would ‘knows’ behave, if it were sensitive to shifting respects of similarity? One would expect that, in contexts where one sufficiently emphasizes *s*’s visual evidence, the nearest world in which *s* does not have hands would be one in which that evidence was held fixed but the reality behind the appearances shifted. That is, one would expect that emphasizing evidence should induce skepticism!

Moving to (D₃), (S) does not fit contextualist solutions to skepticism, for two reasons. First, it is unclear why raising skeptical doubts should generate *any shift at all* in the standards. Why should raising doubts about whether one is a brain-in-a-vat have any impact on the standards at all, much less drive them so far into logical space?¹¹

Second and most crucially, raising doubts would shift the standards in the wrong way for *moderate skepticism*. When standards shift, they do so in ways that *globally infect* other truth-values in that context. Thus suppose again that unresolved doubts are raised as to whether one knows where one is parked. By (S), what must have happened is that the standard *l* has been set to include a relatively wide sphere of worlds (given metric *m*), such that worlds at which *s* cannot track *p* (e.g., car theft worlds) are included. But now, assuming that movie schedule misprint worlds are at least as near as car theft worlds on *m* (an eminently reasonable assumption), it will follow that ‘You know that the movie starts at nine’ counts as false. And assuming that street sign prank worlds are at least as near as car theft worlds on *m*, it will follow that ‘You know that you are now walking on Main’ counts as false, etc. Yet that seems wrong—no *specific doubts* have been raised with respect to these claims. The movie schedule misprint and street sign prank scenarios should not count as relevant, where they have not been raised. Spheres encompass too much.

Shifting finally to (D₄), (S) does not keep score of the question. The globality of standards would conflate distinct lines of inquiry. Thus compare again the inquiries into whether (i) there is a goldfinch in the garden or a blue jay, (ii) there is a goldfinch in the garden or a canary, and (iii) there is a goldfinch in the garden or at the neighbor’s. Whichever scenario (blue jay, canary, neighbor’s) is most distant will subsume the others. The use of a standards parameter will conflate resolving the inquiries associated with the most distant scenario(s) with resolving all the other inquiries associated with equal or nearer scenarios, regardless of any differences in presupposition. (S) thus submerges specific differences in what is posed and what is presupposed, under a general sphere.

Spheres simply are *too topologically limited* to keep score of questions. One might ask whether there is a goldfinch in the garden, or a fiendishly designed robot bird. Or one might ask whether one has hands, the claws of a great

¹¹ DeRose (1995, §12) says that one must track to the nearest *-p* world, and farther to any contextually salient world (his *rule of sensitivity*). This does connect the raising of skeptical doubts with the expansion of the standards. But, as far as I can see, it is invented purely for such purposes, and has no independent support or precedent. It is just a stipulation.

horned owl, the tentacles of an octopus, or the pincers of a crab. Such questions do not describe a sphere. On any reasonable setting of the similarity metric m , such questions denote discrete cells in logical space, with the clumpy topology of cottage cheese (see Figure 5.5). In general, there is nothing in the nature of a question that forces (or even favors) spheres. Thus (S) imposes a topology far too restrictive for inquiry.

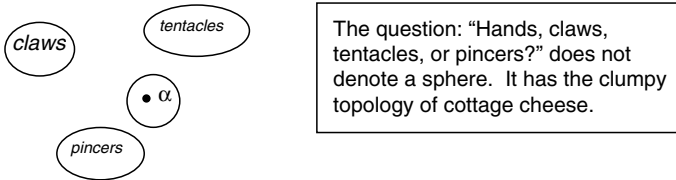


Figure 5.5

The same moral emerges with standards as with thresholds. What shifts with ‘knows’ needs to be *locally responsive to specific doubts and questions*. Standards are too monolithic.

5. Alternatives

I now turn to a third possible answer (rooted in relevant alternative views such as Austin (1946) and Dretske (1981), endorsed by Lewis (1996)), according to which what shifts is the set of *epistemic alternatives* in play. More precisely, for a subject s with a belief p , s is assigned an absolute eliminatory power e for p , where e is the set of possibilities that s can *eliminate*.¹² What shifts is whether e covers enough alternatives for ‘knows’. The way this shifts is that there is a set of relevant alternatives Q , and ‘knows’ requires that all relevant possibilities be eliminated (e must cover Q). Thus:

A. What shifts is the range of alternatives s must eliminate.

Picturesquely, think of s ’s eliminatory power e for p as some region of arbitrary topology in logical space, and think of the relevant alternatives q as some other region also of arbitrary topology. What shifts is whether e covers Q . Thus picture the p -worlds as some box containing actuality, and picture the alternatives as the surrounding boxes, with s ’s eliminatory power e

¹² The notion of elimination is found in Dretske 1981 and Lewis 1996, *inter alia*. According to Lewis 1996, a possibility w is eliminated for s iff s ’s experience in w would differ from actuality. Though other definitions of elimination are, of course, possible.

for p as the shaded boxes (see Figure 5.6). Now suppose that context₁ selects only shaded worlds as relevant alternatives. Then s 's belief that p counts as 'knowledge' in context₁. But suppose that context₂ selects only unshaded worlds, and that context₃ selects a mixture. Then s 's belief that p does not count as 'knowledge' in context₂ or context₃. So far, so good.

But does (A) respect desiderata (D₁)–(D₄)? I believe that the answer is *yes* in each case. Or at least, I believe that the answer is *better than (T) or (S)* in each case.

Starting with (D₁), there are precedents for semantic sensitivity to alternatives, such as with the *modal auxiliaries* (e.g. 'can' and 'must'). Claims like 'wood must burn' and 'I can run a four minute mile' have context-variable truth conditions, depending on which worlds count as *accessible*. 'Wood must burn' is true iff in every accessible possible world, wood burns. Whether this is true depends on whether worlds with different laws of nature are relevant. 'I can run a four minute mile' is true iff in some accessible world, I run a four minute mile. Whether this is true depends on whether worlds in which I am a (vastly) better athlete are relevant.¹³

The modal locutions are a decent precedent for the contextual variability of 'knows', given that 'knows' denotes *epistemic necessity*: K is box.¹⁴ Hence it would be unsurprising if K can shift truth value depending on which alternatives are in play. Box always does. Thus (A) provides a linguistically general parameter, associated with a natural class of expressions of which 'knows' is an instance: the propositional attitude verbs, which model as modalities.

Perhaps an even better precedent is the verb 'regrets'. It is very plausible that the truth condition for regret ascriptions shift with the implicit alterna-




Figure 5.6

¹³ See, for instance, Angelika Kratzer 1977 and 1991 for further discussion of modals.

¹⁴ Or at least, in virtually every epistemic logic, K is treated as box. For some developments of this idea, see Jaakko Hintikka 1962 and G. E. Hughes and M. J. Cresswell 1966.

tives. For instance, 'I regret that Bush is president' is true when the relevant alternative is that Gore is president, but false when the relevant alternative is that Cheney is president. Whether 'I regret that Bush is president' is true seems to depend on who the alternative is.

The verb 'regrets' is an ideal precedent for 'knows', since both are members of the following lexical kind: *factive attitude verbs that permit either declarative (that-clause) or interrogative (wh-clause) complements*. (This kind also includes 'forgets', 'learns', and 'discovers', *inter alia*.) Indeed, given that 'regrets' is semantically alternatives-variable, the default hypothesis ought to be that 'knows' is too.

Turning to (D₂), as mentioned above, what seems to get us to shift between 'knows' and 'does not know' is the invocation and revocation of specific doubts. If the invocation and revocation of specific doubts is understood in terms of the relevance and irrelevance of specific alternatives, then all works. Think of how we may shift from 'must' to 'might not' and back. Or from 'regrets' to 'does not regret' and back. Here we invoke and revoke specific scenarios. (That said, further investigation is called for as to how closely shifts with modal auxiliaries and shifts with knowledge ascriptions parallel. (A) predicts a perfect parallel, once other factors are controlled for.)

Moving to (D₃), what is needed is an explanation for why radical skeptical scenarios generate global doubts, while moderate skeptical scenarios generate local doubts. To begin with, (A) explains why raising skeptical doubts should expand the alternatives. And (A) explains why radical doubts, such as whether one is a brain-in-a-vat, induce global skepticism. For the brain-in-a-vat hypothesis (or some suitable subcase of it) is an alternative to virtually every contingent proposition about the external world. And (A) explains why moderate doubts, such as whether one's car has been stolen and relocated, only induce local skepticism. For the car theft hypothesis (or some suitable subcase of it) is an alternative to the proposition that my car is parked on Elm. But it is not an alternative to the proposition that the movie starts at nine, or that I am now walking on Main, etc.

(T) and (S) had trouble with the local doubts of moderate skepticism, because thresholds *draw lines* and standards *draw spheres*, both of which are globally encompassing. (A) fares better because alternatives themselves are *pointlike* in logical space. That *w* is an alternative has no implications for whether other scenarios (be they associated with comparable degrees, or ringed around the same logical spheres, or whatnot) are relevant. Alternatives are not monolithic.

Shifting finally to (D₄), what is needed is a parameter capable of illuminating inquiry by scoring the question. Alternatives work perfectly, because alternatives are exactly what questions denote. All well-formed questions are multiple-choice questions. As James Higginbotham writes, 'An *abstract question* [is] a nonempty *partition* Π of the possible states of nature into *cells*' (1993: 196) These cells are the semantic image of a (possibly infinite) *multiple-choice slate*.¹⁵

(T) and (S) had trouble with responding to specific questions, and distinguishing what is posed and what is presupposed. (A) fares better because alternatives can describe arbitrary topologies. Thus the question: 'Hands, claws, tentacles, or pincers?' denotes discrete cells in logical space with the topology of cottage cheese. It corresponds to the set of alternatives: {hands, claws, tentacles, pincers}. Thus the inquiries concerning the goldfinch in the garden pose different alternatives: {goldfinch in the garden, blue jay in the garden}, {goldfinch in the garden, canary in the garden}, {goldfinch in the garden, goldfinch at the neighbor's}. And they differ in presupposition, where what is presupposed is simply the union of what is posed.¹⁶ What emerges is that only (A) has the flexibility to handle the full range of questions.

I conclude that what shifts with context are the epistemic alternatives. The epistemic gear that the wheels of context turn is the set of relevant alternatives. Or at least, as between thresholds, standards, and alternatives, with respect to the desiderata (D₁)–(D₄), alternatives are the runaway winner.

6. The Linguistic Basis

What constraints does a parameter of alternatives impose on the linguistic basis for contextualism? I have argued (§5) that an alternatives parameter is linguistically plausible, in so far as it fits the precedents of modal locutions, and of 'regrets'. Thus, the question of the linguistic basis for contextualism becomes the question of the linguistic basis for the accessibility relation for modals, and the alternatives-variability of 'regrets'.

¹⁵ The association of questions with multiple-choice slates is known as *Hamblin's dictum* (C. I. Hamblin 1958), and is implemented in Nuel Belnap and Thomas Steel's (1976) erotetic logic, and maintained in the leading linguistic treatments of interrogatives, such as Jeroen Groenendijk and Martijn Stokhof 1997.

¹⁶ Question *Q* presupposes proposition *p* iff *p* is entailed by all answers to *Q* (Belnap and Steel 1976). Picturesquely, what is entailed by all the cells is that the truth lies in those cells and not outside them.

Here I leave open whether the accessibility relation for modals, and the alternatives parameter for ‘regrets’, should be understood in terms of hidden syntactic variables, semantic indexicals, general indices of semantic evaluation, or whatnot. Though with precedents in hand, the way to approach the issue becomes clear: just figure out how things work in those cases.¹⁷

And with precedents in hand, the worry that there might be *no* linguistic basis for contextualism dissolves. Those who argue that contextualism is linguistically baseless must either (i) argue that ‘regrets’ is not semantically alternatives-variable, despite appearances; or (ii) maintain that ‘regrets’ is not fit precedent for ‘knows’, despite their kinship; or (iii) concede a precedent for a form of contextualism, in which what shifts are the alternatives.

REFERENCES

- Austin, J. L. (1946). ‘Other Minds’, *Proceedings of the Aristotelian Society*, 20 (supp.): 149–87.
- Bach, Kent (*forthcoming*). ‘Context *ex Machina*’, in Z. Szabo (ed.), *Semantics vs Pragmatics* (Oxford: Oxford University Press).
- Belnap, Nuel, and Thomas Steel (1976). *The Logic of Questions and Answers* (New Haven: Yale University Press).
- Castañeda, Hector-Neri (1980). ‘The Theory of Questions, Epistemic Powers, and the Indexical Theory of Knowledge’, in P. French, T. Uehling, Jr., and H. Wettstein (eds.), *Midwest Studies in Philosophy, v. Studies in Epistemology* (Minneapolis: University of Minnesota Press), 193–238.
- Chisholm, Roderick (1988). ‘The Indispensability of Internal Justification’, *Synthese*, 74: 285–96.
- Cohen, Stewart (1988). ‘How to be a Fallibilist’, *Philosophical Perspectives*, 2: 91–123.
- (1999). ‘Contextualism, Skepticism, and the Structure of Reasons’, *Philosophical Perspectives*, 13: 57–89.
- DeRose, Keith (1995). ‘Solving the Skeptical Problem’, *Philosophical Review*, 104: 1–52.
- Dretske, Fred (1981). ‘The Pragmatic Dimension of Knowledge’, *Philosophical Studies*, 40: 363–78.

¹⁷ See Schaffer 2004 and *forthcoming* for some further discussion. See Bach *forthcoming* for further options, including the option that what is said may be *incomplete*, in underdetermining which proposition the speaker is expressing. On this option, knowledge ascriptions may underdetermine the value of the alternatives variable *Q*. But knowledge propositions would invariably have the ternary, relational form: *s* knows that *p* relative to *Q*.

- Feldman, Richard, and Earl Conee (1985). 'Evidentialism', *Philosophical Studies*, 48: 15–34.
- Fine, Kit (1975). 'Vagueness, Truth, and Logic', *Synthese*, 30: 265–300.
- Groenendijk, Jeroen, and Martijn Stokhof (1997). 'Questions', in Johan van Benthem and Alice ter Meulen (eds.), *Handbook of Logic and Language* (Amsterdam: Elsevier Science Publishers), 1055–1124.
- Hamblin, C. L. (1958). 'Questions', *Australasian Journal of Philosophy*, 36: 159–68.
- Heller, Mark (1999). 'Contextualism and Anti-Luck Epistemology', *Philosophical Perspectives*, 13: 115–29.
- Higginbotham, James (1993). 'Interrogatives', in Kenneth Hale and Samuel Jay Keyser (eds.), *The View from Building 20: Essays in Honor of Sylvain Bromberger* (Cambridge, Mass.: MIT Press), 195–228.
- Hintikka, Jaakko (1962). *Knowledge and Belief* (Ithaca, NY: Cornell University Press).
- Hookway, Christopher (1996). 'Questions of Context', *Proceedings of the Aristotelian Society*, 96: 1–16.
- Hughes, G. E., and M. J. Cresswell (1996). *A New Introduction to Modal Logic* (London: Routledge).
- Kennedy, Christopher (1999). *Projecting the Adjective: The Syntax and Semantics of Gradability and Comparison* (New York: Garland).
- Kratzer, Angelika (1977). 'What "Must" and "Can" Must and Can Mean', *Linguistics and Philosophy*, 1: 337–55.
- (1991). 'Modality', in A. Von Stechow and D. Wunderlich (eds.) *Semantics: An International Handbook of Contemporary Research* (Berlin: de Gruyter), 639–50.
- Lewis, David (1979). 'Scorekeeping in a Language Game', *Journal of Philosophical Logic*, 8: 339–59.
- (1996). 'Elusive Knowledge', *Australasian Journal of Philosophy*, 74: 549–67.
- Nozick, Robert (1981). *Philosophical Explanations* (Cambridge, Mass.: Harvard University Press).
- Schaffer, Jonathan (2004). 'From Contextualism to Contrastivism', *Philosophical Studies*, 119: 73–103.
- (*forthcoming*). 'Contrastive Knowledge', T. Grendler and J. Hawthorne (eds.), *Oxford Studies in Epistemology*, 1.
- Stanley, Jason (2004). 'On the Linguistic Basis for Contextualism', *Philosophical Studies*, 119: 119–46.