

Chapter 7

Shifting standards?

Counting on context

Some words are slippery. Every night, the word ‘tomorrow’ slides forward to pick out a different day of the week. ‘Here’ designates a different place depending on where you are standing. ‘I’ stands for someone different depending on who is speaking; and ‘this’ could be anything at all. Words like ‘big’ and ‘small’ are also tricky: a morbidly obese mouse is in some sense big, but in another sense still small. What about the verb ‘to know’? Is it possible that it also shifts around in some interesting way?

What the other words featured in the last paragraph have in common is context-sensitivity. The context in which these words are used plays a role in setting what they stand for. Some words (like ‘I’ and ‘now’) are sensitive to the speaker’s identity and location in time and space. Others (like ‘big’ and ‘tall’) are sensitive to a comparison class: it takes much more height to be tall for a skyscraper than it does to be tall for a cereal box. To complicate matters, the same thing could be a member of two different classes (mouse, animal): a creature could be big for a mouse and still small for an animal, and which adjective we should use for it depends on which of these classes we have in mind.

It's tempting to say that context-sensitive words keep changing their meaning, but that's not exactly right. We don't have to buy a new dictionary every day to keep up on what the word 'tomorrow' means. There are some fixed rules: 'I' always picks out the person speaking and the relevant sense of 'tall' always means 'large for its kind in the vertical dimension'. Rather than changing their meanings, context-sensitive words work like recipes that take input from a conversational context to settle what they stand for. Once the context is established, it should be clear exactly what 'this' indicates, or which day of the week is picked out by 'yesterday'.

'Contextualism' is the standard name for the view that words like 'know' and 'realize' are context-sensitive. Some of the appeal of contextualism comes from its promise to reconcile the main points in Chapters 1 and 2 of this book. Chapter 1 observed that the verb 'to know' is one of our most common verbs, and is used as the default label for ordinary cases of seeing, hearing, or remembering that something is the case. (Of course, you know that you are reading a book right now.) Chapter 2 observed that it's easy to discover ourselves doubting that knowledge is ever humanly possible. (How could you ever really know that you are reading a book and not just dreaming that this is so?) According to contextualists, there's no real clash between the positive remarks in the first chapter and the negative remarks in the second. Because 'know' is context-sensitive, everyday claims about how much you know are fully compatible with sceptical claims about your knowing almost nothing. The everyday speaker and the sceptic are in different conversational contexts, and are therefore saying different things when they use the word 'know'. Just as the sentence 'Tomorrow is Friday' sometimes says something true, and sometimes says something false (depending on when you say it), knowledge-ascribing sentences like 'John Doe knows that he is reading a book' do a similar trick. When everyday low standards are in force, it's right to say 'John Doe knows'; when we are using the sceptic's high standards, the thing to say is 'John Doe does not know.'

The emergence of contextualism

Contextualism grew out of a theory of knowledge launched in the early 1970s, the ‘Relevant Alternatives’ theory of knowledge. Advocates of that theory say that knowing always involves grasping some kind of contrast. Here’s an example: Jane Roe is at the zoo with her son, and sees a black-and-white striped animal in the enclosure ahead of her. ‘Look, Billy!’ she says. ‘That’s a zebra!’ She’s right; there’s nothing wrong with Jane’s eyesight or her ability to recognize ordinary zoo animals, and the black-and-white striped animal she is looking at is indeed a zebra. When asked whether Jane Roe knows that the animal she is looking at is a zebra, we find it easy to say ‘yes’. But here’s a trickier question: does she know that the animal she is looking at is not a *cleverly disguised donkey*? (Figure 8). It’s possible to paint stripes on a donkey and trim its ears and tail, and from where Jane is standing, a cleverly disguised donkey would look exactly the same to her. Jane Roe’s evidence, according to the Relevant Alternatives theory, is good enough to enable her to tell that the animal is a zebra (as opposed to a lion, antelope, or camel). Given the range of animals at an ordinary zoo, she has a pretty easy set of relevant alternatives to pick from. However, her evidence is not good enough to enable her to tell that the animal is not a cleverly disguised donkey: to make that harder judgement, she’d have to rule out the relevant alternative that it actually is a cleverly disguised donkey. It is not impossible for her to rule out that tricky alternative, but Jane would need to hop the fence and go closer, perhaps even close enough to dab some cleansing fluid on the animal’s fur. So, from 20 paces away, she knows that the animal is a zebra, but she doesn’t know that the animal is not a cleverly disguised donkey.

In this example, it’s understood that being a cleverly disguised donkey is not a relevant alternative at an ordinary zoo like the one Jane Roe is visiting, but it could be a relevant alternative at an impoverished zoo, or at a zoo run by practical jokers. One of the challenges facing the Relevant Alternatives theory is coming up

with the rules for setting what the relevant alternatives are for any given claim. The field of alternatives is partly set by what is actually possible for the speaker, but the further details of the theory have been hard to work out. Does it matter whether Jane Roe knows that the zoo is not run by jokers? It's unclear. But more importantly, there's something deeply *strange* about the Relevant Alternatives theory.



Figure 8. A man stands with one of Marah Zoo's world famous painted donkeys, in Gaza City, Palestinian Territories, December 2009

Let's assume that Jane knows some basic biology: suppose she knows, as most adults do, that zebras and donkeys are different species of animal, so that nothing can be both a zebra and a donkey at the same time. If the animal is a zebra, then it follows

from basic biology and logic that the animal is not a donkey. Logic also tells us that if something is not a donkey, it is not a cleverly disguised donkey. What the Relevant Alternatives theory says is that as Jane stands 20 paces away from the zebra, she knows that the animal she sees is a zebra, and also knows that if it is a zebra it is not a cleverly disguised donkey, but she doesn't know that the animal is not a cleverly disguised donkey. She can know the premises of the simple argument in Box 6 without being able to know the conclusion, despite being able to see for herself that it follows logically. The Relevant Alternatives theory therefore violates a principle called Closure, according to which you know anything that you succeed in deducing logically from your existing knowledge.

Box 6 A problem for the Relevant Alternatives theory of knowledge

What Jane knows from 20 paces away:

- That animal is a zebra.
- If that animal is a zebra, it is not a cleverly disguised donkey.

What she doesn't know:

- That animal is not a cleverly disguised donkey.

Violations of Closure seem weird: can't you trust logical deduction? Furthermore, as soon as we say that Jane doesn't know that the animal is not a cleverly disguised donkey, it sounds odd to insist she still knows it's a zebra. However, there is something compelling about the initial observations that motivated the Relevant Alternatives theory: it sounded right at first to describe Jane as knowing that the animal was a zebra, and it also sounded

right to deny that she knew at a distance that the animal was not a donkey in disguise.

Contextualism emerged as a way of keeping the appealing part of the Relevant Alternatives Theory without accepting its strange denial of Closure. Published in 1976 by Gail Stine, the first clear formulation of contextualism was a two-part proposal. Stine's first claim was that we use higher or lower standards for knowledge in different settings: 'It is an essential characteristic of our concept of knowledge that tighter criteria are appropriate in different contexts. It is one thing in a street encounter, another in a classroom, another in a law court—and who is to say it cannot be another in a philosophical discussion?' Rather than having a field of alternatives fixed by what is possible for the person who is making the judgement, Stine proposed that various narrower or wider fields of alternatives are surveyed by those who are talking about whether a given person has knowledge. Stine's second point was to insist that within any given context, we appropriately stick to one set of standards. In a context in which we are worried about the possibility of doctored donkeys wearing make-up, it is wrong to say, 'Jane knows that the animal she is looking at is a zebra.' If she doesn't share our worries, Jane can truly say, 'I know that's a zebra,' but she can't then use pure logic and biology to deduce the exotic conclusion that it's not a cleverly disguised donkey. As soon as Jane starts thinking about cleverly disguised donkeys, her standards for knowledge will rise to take such exotic possibilities into account, and it will become wrong for her to say 'I know it's a zebra.' It's legitimate to talk about knowing with high standards or low standards; you just can't slide back and forth between the two in a single context without marking the shift.

In Stine's formulation of contextualism, knowing requires discrimination from a larger or smaller field of alternatives. The everyday speaker wonders whether the sandwich is chicken or tuna, and counts a person as knowing if he can rule one of those options out in favour of the other; the sceptic wonders about many more exotic alternatives as well (hypothetical new kinds of tuna that look like chicken, mere holograms of sandwiches, projections

sent by an evil demon, and so forth). Once we have this wider array of alternatives to contend with, it's harder to credit anyone with knowledge. The widening and narrowing of a field of relevant alternatives is one vivid way in which contextualism can be expressed. But the basic idea behind contextualism is something more general, and not necessarily married to the relevant alternatives idea: the idea is that 'know' expresses something different as situations change, and different contextualist theories have different lines about how this works. There are internalist formulations of contextualism in which different contexts call for more or less evidence, and externalist formulations of contextualism in which belief-forming processes must track the truth across a narrower or broader array of circumstances as standards rise or fall. Contextualism on its own is not a theory of knowledge: it's a theory about knowledge-attributing language, a semantic topping that can be spread onto various different underlying theories of knowledge.

Is it an appetizing topping? Contextualism does promise a neat solution to the problem of scepticism: both sides are right, in a sense. The man in the street who says, 'I know that I am reading a book' is saying something true, but so is the sceptic who says, 'The man in the street does not know that he is holding a book.' The trick is that the word 'know' picks out a different relation for each of these speakers. Talking about scepticism raises the standards for knowledge, just as talking about basketball players raises the standards for 'tall', to use a favourite contextualist analogy. A man who measures six feet (183 cm) ordinarily counts as tall in the United States of America, where the average male height is about 5'9" (175 cm), even if six feet doesn't count as tall in the National Basketball Association, where the average height is about 6'7" (201 cm). Consider Chris Paul, a six-foot player for the LA Clippers. Sports fans discussing the strengths and weaknesses of the team can honestly say 'Chris Paul isn't tall' at the very moment that Chris Paul is honestly describing himself as tall on his dating website profile. There is no real conflict here, thanks to the difference in standards.

We have to be careful about speakers who are talking about what other speakers say: it would be a mistake for Chris to say ‘What those sports fans said about me was false’ simply on the basis of his being above average in height for an American man; it would also be a mistake for the fans to accuse Chris of saying something false on his dating profile simply on the basis of his being short for a professional basketball player. When you judge the truth or falsity of other speakers’ context-sensitive remarks, you need to respect their context (and they need to do the same for you). The most straightforward thing to do would be to make the comparison class explicit. The sentence ‘Chris Paul is tall’ leaves a blank (compared to what?) that has to be filled in by context; both of the sentences ‘Chris Paul is tall for an American man’ and ‘Chris Paul is not tall for a professional basketball player’ have that blank filled in, and so express truths in a stable fashion across both contexts.

Going back to the case of knowledge, what’s the clearest thing to say? Can we make our standards explicit? Is ‘You know that you are reading a book’ just like ‘Chris Paul is tall’? If so, we should just be able to solve the problem of scepticism by saying something like ‘You know, by low standards, that you are reading a book, but you don’t know, by high standards, that you are reading a book.’ Well, that was easy. But was it satisfying?

It’s hard to say. Defending the context-sensitivity of ‘know’ isn’t quite as smooth as defending the context-sensitivity of ‘tall’. One objection that is often raised against contextualism is that it is really just scepticism in disguise. A sceptic doesn’t have to fight the low-standards notion of knowledge in play in ordinary life; sceptics themselves doubtless use the verb ‘know’ dozens of times every day when they are talking to each other about ordinary things. What the sceptic really wants to say is that when we look closely and carefully at knowledge, we’ll see that our everyday claims to have it are actually false. This doesn’t amount to a call for stamping out all talk of knowledge, because we may succeed in communicating something useful in our casual talk about knowing. It could be interesting to hear a co-worker say, ‘Lee knows who

got the job,' for example. Thinking about the strict meaning of 'know', we may decide that Lee doesn't really know who got the job (could Lee really guarantee that he is right, and that the apparently successful candidate hasn't just been struck by a meteorite?). Maybe by saying 'Lee knows,' what's really communicated is that k's a good bet Lee could tell you. It's not crazy to think that you can convey something useful by saying something literally false: you use many words in this loose manner, like saying 'I'm starving,' when you are really just hungry. Even if that's literally false, it can help to get your host to offer you a snack. In the sceptic's high-standards context, 'Lee knows who got the job' expresses something false, even if we do something useful by saying that sentence in everyday contexts. But if contextualism says that sceptics speak truly when they deny that ordinary people know simple facts (like the fact that you are reading this book), then contextualism seems to be letting scepticism win, at the expense of common sense.

Contextualism is actually more subtle than the rough 'loose use' theory sketched above. In particular, contextualists are careful about respecting the context of other speakers. According to contextualism, sceptics get to say, truly, 'Lee doesn't know who got the job.' But they *don't* get to say that everyday speakers are saying something false when those everyday speakers say, 'Lee knows who got the job' in an everyday context. The same reasoning applies on the other side: ordinary people can express something true by saying, 'I know that I am reading a book,' in an ordinary context, but they cannot say that the sceptics are speaking falsely in saying, 'You don't know that you are reading a book' in the context of a high-standards philosophical discussion. Contextualists themselves don't have to say that the sceptical philosopher's context is better (or worse) than the common person's: high standards are not necessarily better than low standards (in fact, high standards can be really annoying and pedantic when you just want to send a card congratulating the person who got the job). Both sides win, as long as they play nice and refrain from putting down what the other side is saying.

Contextualism's tender tolerance for other points of view does not appeal to everyone. Critics of contextualism continue to resist the idea that the sceptic and the common man are both saying something true, and continue to wonder which way of talking *really* manages to get it right about knowledge itself, once and for all. From a contextualist perspective, asking which way of talking ultimately captures the nature of knowledge is like asking which weekday is ultimately 'tomorrow'. It's not a good question.

If contextualism aims to be somewhat friendly to both the sceptic and the common person, it has to be less friendly to the philosophical adversary who thinks there is a single true answer to the question 'Does Lee know who will get the job?' no matter which context we ask it in. That person really is making a mistake, contextualists will say, and they will say the same about anyone who thinks that the common person and the sceptic can't both speak truly. Contextualists do recognize that it's very common to think that one is forced to take sides: somehow the way in which 'know' is supposed to shift is more hidden from us than the way in which 'tall' or 'here' is. We don't have a big philosophical tradition of debates about which place is *really* 'here', in the way that we have debates about which way of talking is really right about knowledge. But why exactly would the context-sensitive workings of our language be obscured to us when we are talking about knowledge, if they are so transparent when we are talking about times, places, and qualities like 'tall'? One of the most active current research questions for contextualists concerns just this question, and various proposals have been advanced. Perhaps something about our use of 'know', like its role in closing off further enquiry, hampers us from tracking context shifts as well as we should, and gives us an illusion that knowledge is absolute. Or perhaps contextualism is wrong, and knowledge itself really is absolute. The view that knowledge is absolute, in the sense that the words we use for it are not context-sensitive, is known as 'invariantism'. Invariantism faces a challenge in explaining the

shifting intuitions that make knowledge sometimes seem easy and sometimes seem hard.

Interest-relative invariantism

Lee is walking towards the bus stop after work when he bumps into his co-worker Smith, who is headed back into the building.

- Lee, do you know if the ground-floor supply room door is locked? I just realized I left my jacket in there, and I don't have a key.'
- 'Yes, it is—I know because I locked it myself half an hour ago, and I didn't see anyone else in that hallway after. Sorry!'

Lee doesn't seem to be saying anything controversial in claiming knowledge here, and we'll assume as we go forward that in fact the door really is locked, and that Lee's key, his eyesight, and his memory are all fine. But now we'll imagine a different version of the story, one that takes a different turn after Lee leaves the building. As Lee walks towards the bus stop, he is approached not by Smith, but by four police officers.

- 'Excuse me, sir, but we have an emergency situation in the building you just left. Apparently there was a shooting on the second floor, and the gunman is still in the building. Are there any ways out of the building other than the front door here?'
- There's a back door leading off the supply room, but I locked the door to that room half an hour ago.'
- Do you know if it's still locked, or if anyone else might have opened it?'
- I don't know—I didn't notice anyone go by, but I wasn't watching the door the whole time.'

In the forgotten-jacket version of the story, Lee claimed to know that the door was locked; in the gunman version, he claimed not to know. Both times he was saying something that sounded true. The curious thing is that both stories ran parallel up to the moment Lee left the

building: in both stories he is trusting his memory of the last half hour as he answers the question about whether he knows. In traditional epistemology, whether or not you know depends on traditional factors such as whether your belief is true and how good your evidence is: interestingly enough, all these factors seem to be the same in both stories. So how is it possible that Lee knows the door is locked in the first, but doesn't know in the second?

It's clear what contextualists would say: in the casual bus-stop conversation, there are low standards in play, and when the police get involved, the standards rise. Lee is saying something true when he says 'I know' in the first story, but also when says 'I don't know' in the second. But this is not actually a story about how Lee knows in the first story and not in the second: it's a story about what a person can say truly in the two contexts. Viewed from other ('higher') perspectives, it would be *false* to say that Lee knew the door was locked in the first story, or so contextualism maintains. There is no simple, context-independent answer to the question 'Did Lee know or didn't he?'

If we were pushed towards contextualism by the feeling that it seems like Lee really does know in the first story, and doesn't in the second, then we might try to find a theory of knowledge which would fit those feelings more directly. If all the traditional factors that matter to epistemology (truth, evidence, reliability, and so forth) are the same in the two stories, one possibility is to allow some non-traditional factors to make a difference. What other factors are different in the two stories? Advocates of a position now known as interest-relative invariantism (IRI) have noticed that there are *practical* differences between the two stories. In the first, not much is at stake for Lee; as a friendly colleague he'll go back and unlock the door so Smith can get his jacket, but if he's wrong about the door still being locked, it's not a big problem. Finding it open would mean only that he'd wasted a minute on a short walk. In the second story, there could be serious practical consequences if Lee is wrong about the door being locked: the gunman could escape out the back exit.

Interest-relative invariantists have found that practical interests seem to have an impact in many cases. Do you know whether the sandwich is chicken or tuna? If not much is at stake (you like chicken just a little bit more, and you would slightly prefer to get soup instead if it's tuna), then a fairly casual inspection (looks like chicken) would be enough to count as knowing. If your life is at stake (you have a very serious fish allergy), then you will not know on the basis of that casual inspection, according to IRI. The more you have at stake, the more evidence you need in order to count as knowing. This is how the two versions of the Lee story can deliver different verdicts about whether Lee knows.

What makes IRI different from contextualism is that IRI is a theory about how knowledge itself works, not just a theory about the semantics of knowledge-ascribing vocabulary. The verdicts that it delivers are not just true-when-expressed-in-certain-contexts, but true full stop Lee knows in the first story; he doesn't know in the second. The context that matters in setting the standards for how much evidence Lee needs for knowledge is Lee's own context, not the contexts of other people who might be talking about him from different perspectives. According to IRI, the sceptic is just wrong to say that Lee lacks knowledge in the first story. Different advocates of IRI have different accounts of the nature of knowledge; what they all have in common is that practical interests, which are not a factor for traditional epistemologists, are a factor that helps determine whether or not a person knows: as stakes rise, more evidence is needed for knowledge. Advocates of the theory contend that IRI is the best way of making sense of the relationship between knowledge and action.

Old-fashioned invariantism, again?

Contextualists were quick to criticize the IRI approach, noticing, for example, that it also has trouble capturing some patterns of shifting intuitions. If it's a plain context-independent fact that Lee knows that the door is locked in the first (forgotten-jacket) version

of the story, how is it that the sceptic is able to make us start doubting that fact so easily? The IRI approach also gets awkward when we talk about counterfactual possibilities. ‘The waiter doesn’t know whether the sandwiches are tuna or chicken, but he would have known if it weren’t for the fact that one of his customers has an allergy.’ That sounds odd, but IRI predicts it shouldn’t. Meanwhile, advocates of IRI have fired back at the contextualists, often by pointing out that ‘know’ doesn’t really work like other context-sensitive vocabulary: unlike ‘tall’, it doesn’t easily fit on a sliding scale, and unlike ‘today’, there are no simple rules for explaining how context fills it in. Contextualists have suggested that ‘know’ might have a special kind of context sensitivity all its own, but it’s still an open question exactly how it works.

As the main shifting-standard views take shots at each other, advocates of more rigid standards have wondered whether their old-fashioned view might still win the day. Strict invariantists maintain that knowledge is determined strictly by traditional factors (truth, evidence, and the like) and that knowledge-attributing vocabulary is not sensitive to context. We’ve already met one of the most straightforward forms of strict invariantism: scepticism. According to Academic Scepticism, for example, there is a single fixed standard that must be met in order to have knowledge (we must have an infallibly correct impression of the thing judged). Sadly, we never meet this standard in daily life (or perhaps we meet it only for one or two special claims, like ‘I exist’). If you take yourself to know you are reading a book, you are just wrong. Sceptics seem to owe us a story about why we speak so much of knowledge if it is forever out of reach, and we may or may not be satisfied with what they have to say. If we are unsatisfied, we may want to turn in the direction of moderate strict invariantism, which holds that there is a single fixed standard that must be met in order to have knowledge, but also that it is a standard that humans often meet. You do know that you are reading a book, and Lee really did know that the door was locked. Strict moderate invariantists also owe us something: they need to explain why exactly the sceptic can so easily lead us to doubt our

everyday judgements, and why Lee's knowledge seemed to melt away as he was questioned by the police, even though the strictly traditional factors he was relying on remained the same. Strict moderate invariantists have struggled to answer these questions. One avenue they have tried is to argue that there is something wrong with our shifting intuitions, or with the cases that produce them. Perhaps the differences between the two versions of the locked-door story are larger than they seem: we assumed that the traditional factors that matter to knowledge were identical across those cases, but it's possible that the high stakes in one case will naturally trigger lower confidence or a different way of thinking about one's evidence. Or perhaps something about the high-stakes situation makes us confused about the difference between knowing and knowing that we know, or about the difference between what we are literally saying and what we are trying to convey. Perhaps there is something wrong with our instincts about these cases; perhaps some natural distortion is introduced when we talk to the sceptic or weigh life-and-death issues. Given how hard it is to develop a smooth story about our patterns of intuition about knowing, it makes sense to take a deeper look at how those intuitions are produced.