

The Attitude of Knowledge

If John Doe wants to know whether his bank will be open this coming Saturday, he could try to remember whether he's been there on a Saturday in recent memory. He could make a phone call. Or he could walk to the bank, look at the sign in the window, even interview the branch manager. Moving further along what might be an open-ended spectrum, he could collect data on the reliability of posted signs, the accuracy of branch managers' statements concerning ordinary operating procedures, and the rarity of unexpected branch closings due to events such as fires.

But precisely when can one say that John Doe *knows* his bank will be open? DeRose has an interesting story about why this question is hard to answer. His well-crafted examples strongly suggest that different circumstances can make different levels of evidence seem adequate. DeRose takes his examples to show that 'know(s)' has variable semantic content, where the part that is variable concerns something he dubs 'strength of epistemic position.' Other things being equal, our John Doe will be in a weaker epistemic position if he just consults his memory, and a stronger epistemic position if he also talks to the manager, but there is no single right answer to the question of just how strong an epistemic position he needs to be labeled a knower. Ascribers with different conversational purposes will impose different epistemic standards, different thresholds for how far along the spectrum they require a subject to venture in order to count as knowing, and the semantic content of 'know(s)' will vary accordingly.

On this theory, the semantic value of an utterance of 'know(s)' is set by the epistemic standards of the ascriber, which are in turn a function of her conversational

purposes. Consistent success in figuring out whether utterances involving that verb express truths would therefore require some systematic understanding of both the conversational purposes normally elicited by various situations, and also the types of standard appropriate to those purposes. For some context-sensitive terms (like 'I' and 'now'), a simple rule governs the relationship between the context of utterance and the semantic value of the term; it seems the corresponding rule for 'know(s)' would be considerably harder to state. DeRose himself does not aim to state the rule (or explain how it would be learned), but its complexity becomes evident when he describes the natural standards for his core examples. According to DeRose, speakers naturally employ low standards when they aim to predict or explain the behavior of those facing low-stakes decisions, and high standards when tracking high-stakes. Meanwhile high-stakes speakers naturally employ high standards when considering third parties who are being considered as sources of testimony, even if the third parties themselves are in low-stakes situations. DeRose thinks ordinary people grasp these relationships between knowledge-ascribing speakers, subjects and standards, and can tell whether ordinary 'knowledge'-ascribing sentences express true propositions.

I take DeRose's examples to succeed in showing that we have naturally shifting expectations about evidence collection; however, I would not conclude that these shifting expectations are a sign of shifting standards. There is a simple reason why high-stakes situations tend to trigger the expectation of greater evidence collection: evidence-gathering behavior, like any rational behavior, balances expected costs and rewards. Normal cognition is said to be roughly 'adaptive': when stakes are high, additional effort spent searching for evidence is offset by additional accuracy in judgment. However, the fact that

evidence collection is normally adaptive across various circumstances provides a way of reinterpreting the core intuitions motivating contextualism. Rather than taking epistemic standards to shift from one context to another, we might expect different epistemic behavior in different circumstances exactly because of an invariant expectation that normal subjects will think adaptively.

In reading epistemic scenarios it is easy to focus on explicitly mentioned facts about evidence collection; however, implicitly computed representations of the subject's attitude also matter to knowledge, and where variations in evidence collection are naturally driven by variations in attitude, we need to keep track of both sides of the equation.

DeRose himself takes psychological attitude to be a central and invariant requirement on knowledge. From the outset, DeRose maintains that knowledge requires 'roughly, that S has a true belief that p and is in a *good enough* epistemic position with respect to p' (3), and is clear that the only variable factor in this formula is 'how good is good enough' (ibid.). The 'roughly' marks some qualms – to be discussed – about whether 'belief' is the right label for the psychological attitude underpinning knowledge, but DeRose himself allows that outright belief is close enough for most purposes, and this attitude will serve us here until further notice. The rest of DeRose's analysis, meanwhile, is upheld consistently, and exhaustively, in the sense that no conditions beyond these three ever play a role in determining the truth of knowledge-ascribing sentences. The notion of 'strength of epistemic position' is then further constrained: it is strictly a function of factors that are 'truth-relevant', DeRose contends, 'in the sense that they affect how likely it is that the belief is true, either from the point of view of the subject or from a more objective vantage

point' (24). The shift between lower and higher epistemic standards must therefore be understood as a shift between demanding lower and higher reliability on the part of the subject, on some broad conception of reliability that may incorporate subjective and objective factors.

DeRose gives us a series of cases intended to show variation in standards as a matter of sheer rising and falling reliability. These examples involve the collection of more or less evidence: a subject either remembers a visit to the bank, or goes in to check the hours, others judge that a co-worker is present by more or less direct evidence. It is not obvious that we intuitively register differences in evidence as sheer differences in reliability: variations in evidence collection also have an impact on whether a person has an outright belief at all on a given question. Empirical work on judgment shows that no fixed level of evidence produces outright belief across the types of conditions that contextualism has focused on in its discussion of changing standards. We naturally seek more evidence before making up our minds when stakes are high, or when we feel self-conscious, or anticipate having to justify ourselves to others (Lerner & Tetlock, 1999; Payne, Bettman, & Johnson, 1993). If the fixation of outright belief normally terminates the process of seeking and assessing evidence, it is not surprising that typical high-stakes subjects will achieve outright belief only when they have acquired more evidence than their low-stakes counterparts, and weighed it more carefully.

There are conditions under which high-stakes subjects can attain outright belief given only a level of evidence typically just sufficient for their low-stakes counterparts; however, in order to curtail high-stakers' normal rational pursuit of greater evidence, one

needs to compromise their rationality in some manner, for example by putting them under an condition like distraction or time pressure (Kruglanski & Webster, 1996). Because these 'hot' conditions lower accuracy as they raise confidence, it will not in general be correct to ascribe equal reliability to subjects who have the same evidence and level of outright belief but differ in their stakes.

In summary, DeRose is right that when we are evaluating high stakes subjects, we will expect them to engage in more evidence collection than their low-stakes counterparts. But it is not clear that we do so because we expect them to meet a different standard. When we hold the level of evidence fixed, a change in perceived stakes will naturally shift the likelihood that a subject will possess an outright belief. Holding both evidence and level of outright belief fixed, a change in stakes will naturally indicate a significant qualitative difference in the type of basing that underpins the belief, making the belief one of a type that is less likely to be true. So the low-stakes subject who relies on memory about banking hours and the high-stakes subject who also stops in at the branch to check may be engaging in different behavior just to achieve the same standard: each is settling an open question by gathering sufficient evidence for evidence-based outright belief. Ordinary high- and low-stakes subjects may even come across as similarly reliable, if 'reliability' is taken to cover not just expected accuracy with respect to the particular proposition being judged, but more broadly, expected accuracy of the larger cognitive tendency displayed. That tendency – forming beliefs based on the 'cool' basis of reasonably collecting and weighing evidence, whether stakes are high or low – would contrast favorably with the rival 'hot' processes of subjects unable or unwilling to engage in the ordinary rational pursuit of evidence.

We are now in a position to reanalyze DeRose's cases. DeRose sees them as presenting us with superficially contrasting but equally acceptable epistemic evaluations of 'epistemic twins' – parallel subjects who are supposed to enjoy epistemic positions of equal strength and the same psychological attitude to a true proposition. The invariantist aims to show that these parallel characters are not naturally seen as twins in the first place.

This invariantist strategy can be pursued for each of the several psychologically realistic ways of construing the Bank cases. On the reading I find easiest, HIGH can be seen as having suspended his outright belief on the issue of the next day's banking hours at the moment he declares himself not to know. Apparently wanting more evidence than a memory of a recent visit, HIGH's wife has re-opened the question of whether the bank might be closed due to a change in hours, putting HIGH back into the evidence-seeking phase of needing to make up his mind. But if HIGH has temporarily lost his attitude of outright belief, it is easy to see why he lacks knowledge.

DeRose takes such readings of the case to be precluded by his stipulation that in the second case HIGH remains 'as confident as [he] was before that the bank will be open.' (2) If we read this stipulation as a claim that HIGH never loses his outright belief in the face of his wife's doubts, it is not in fact an open question for HIGH whether the bank will be open the next day, even as he declares his intention to go in and "make sure". On this reading he is placating his wife, and only pretending to feel the need to check the hours. Now we face a choice. If we see the wife's request for additional evidence as rational, HIGH's refusal to take it seriously points to some epistemically problematic disposition in him. We may not be sure what motivates his reluctance to rethink the question of the bank's opening in the

light of further evidence – is he strangely dogmatic? – but HIGH’s tendency to answer this question in a cool, evidence-based way is in doubt, and it is no longer clear that he is the reliable epistemic ‘twin’ of LOW. On the other hand, if we do not see the wife’s request for more evidence as rational, then we will see HIGH as reasonable in maintaining his outright belief in the face of that request; we can also see HIGH as maintaining his knowledge throughout. But on this construal the sincerity of HIGH’s denial of knowledge is in question: he knows, we know he knows, and we naturally see him as saying he doesn’t know and going through the motions of checking to please his anxious wife. He may come across as speaking correctly if we naturally perceive it as rational of him to placate his wife in this manner, but what he is saying registers as appropriate rather than literally true. Neither of these interpretations gives DeRose what he needs:¹ a clear example of epistemic twins receiving different but correct epistemic evaluations at the hands of ascribers with different epistemic standards.

There are alternatives to reading the stipulation of equal confidence as a stipulation of continued outright belief. Later in the book DeRose does more to articulate what he means by ‘confidence’, drawing a contrast between a ‘stable’ and an ‘unstable’ sense of the term, and endorsing the former as a better basis for intellectualism. In the ‘unstable’ sense, DeRose grants that HIGH will naturally be seen as less confident after his wife raises a worry about the change in hours: after all, now HIGH is no longer willing to assert the proposition that the bank will be open, and feels the need to make sure. If confidence is

¹ I have not presented an exhaustive list of ways to read the case; others might also be considered. For example, perhaps HIGH knows, but does not know that he knows; on this reading we might see his wife’s demand as rational if iterations of knowledge are required in high-stakes circumstances (See Williamson, 2005 on this point).

taken to be a matter of these 'local' dispositions of action and assertion that we exhibit in particular high- and low-stakes situations, DeRose acknowledges that confidence will naturally sink as our perceived stakes change. However, DeRose thinks that 'confidence' can be understood in a more stable way, as consisting in 'dispositions that tend to remain steady as a subject moves back and forth between high- and low-stakes situations.

Whether a subject is in a high- or a low-stakes situation, she will typically have the same dispositions concerning whether she will flat-out assert a proposition should it happen that it is conversationally relevant and she is in a high-stakes situation.' (2009, 191-2) On this reading, DeRose claims, 'if I'm a fairly normal person, I'm about equally confident in the truth of the relevant proposition in LOW and in HIGH, but the same level of confidence that produces confident behavior in LOW doesn't produce similarly confident behavior in HIGH, because that same level of confidence is not confidence enough to yield that behavior, given the gravity of the situation in HIGH.'

It's an empirical question what is unchanged in 'fairly normal people' as stakes shift. But even accepting the (perhaps controversial) supposition that there is something steady that deserves the name of 'confidence' here, we still face a further question about the relationship between stable confidence and outright belief. DeRose does not explain just what makes a certain fixed level of stable confidence issue in different behavior across LOW and HIGH, but it is open to the invariantist to suggest that this relationship would be mediated precisely by the presence or absence of outright belief (Weatherson, 2005). This suggestion may end up being ruled out by some more detailed theory of stable confidence that DeRose has in mind. I expect DeRose has more to say here: mentioning that he has chosen to bracket this issue for the purposes of this volume, he nevertheless notes that he

finds it 'highly questionable' that plain outright belief is enough to satisfy the psychological attitude requirement of knowledge. He suggests *S is certain that p* might be closer to the mark, and then ventures the more radical idea that 'Perhaps there is no common phrase in natural English that perfectly captures this requirement, in which case we may have to make do with calling it something like 'the attitude of knowledge'. (2009, 186 n1)

However the 'attitude of knowledge' comes to be fleshed out, there are limits to the work that can be assigned to it within the bounds of DeRose's larger ambitions. Because psychological notions like confidence and the attitude of knowledge are on the invariant side of the ledger, it would be problematic for DeRose to assign them an explanatory role which would duplicate the work he has credited to rising and falling epistemic standards. For example, if the invariant psychological factor of stable confidence already determines that our subject with moderate evidence will stop asserting that *p* as stakes rise, it becomes redundant to appeal to the knowledge account of assertion plus rising standards on knowledge to explain this behavior. DeRose risks weakening his case if he can no longer use changing patterns of assertion as evidence of rising standards for knowledge.

Whatever we say about confidence, the line of criticism pursued so far may seem to leave untouched the strongest part of the case for contextualism. Worries about variation in (unstable) confidence or outright belief are easiest to press against first-person cases that identify the subject and ascriber of knowledge. DeRose himself notes that HIGH's worried state of mind might be taken as a sign that HIGH and LOW do not hold positions of equal epistemic strength (2009, 60); he then shifts to third-person cases he takes to be immune to such problems.

In the third-person Office case, two ascribers – Thelma and Louise – know that they share with each other and with a third party, Lena, precisely the same body of evidence for the proposition that John was at work today: all three of them saw his hat in the hall and overheard a partial conversation. Later Thelma and Louise separate and each finds herself discussing Lena’s state of mind, with Thelma in a low-stakes conversation in a tavern saying that Lena knows John was at work, and Louise in a high-stakes interview with the police saying that Lena does not know that John was at work. DeRose claims that we will see Louise as speaking truly when she says that Lena does not know, and that we will do so because we see that Louise’s conversational purposes in speaking to the police would make it natural for Louise to employ high standards. But one wonders whether there might not be a simpler explanation of why Louise’s claim seems true: when the police ask her about Lena’s epistemic position, Louise gives an evaluation of the epistemic position that she naturally expects the police would find in Lena, were they to interview Lena. Louise has found that under pressure to testify, her own evidence is insufficient to settle the question of John’s whereabouts; she naturally anticipates that Lena would have the same experience. The circumstances make it equally natural for Louise to imagine Lena under pressure, and for us, as readers of the case, to imagine Lena the same way, as we perform the second-order mental state ascription task of thinking about what Louise is thinking about Lena. Some invariantists might say that Louise’s assessment of Lena’s current situation (and, following in Louise’s footsteps, our own assessment) is naturally distorted by this imaginative exercise. Others might say that in saying Lena does not know, Louise is just giving the police what they are really interested in (and perhaps what we are also interested in, as readers of the story) – an accurate assessment of the position Lena will be

in when the police encounter her, rather than an assessment of how things are with Lena right now at home. But either way, this looks like a case in which the Lena who is being assessed by high-stakes Louise is not naturally seen as the epistemic twin of the Lena who is being assessed by low-stakes Thelma.

On an ‘adaptive invariantist’ line of thinking, our stakes shift our psychological attitudes, and consequently our evidence-gathering behavior; when we assess others we anticipate that shifting stakes will produce similar effects in them. When we are evaluating others as possible informants, for example, it will be natural for us to imagine them in cooperative conversation with us, obliged to take our high-stakes worries seriously. When we imagine them as actors, and attempt to predict or explain their behavior, we focus on the practical environment in which they find themselves, and expect them to think adaptively in that environment. Although this imaginative representation generally does a good job of predicting how others will think and act, it is subject to certain known hazards: for example, where we have worries not shared by the subject we are evaluating (say, worries about tricky lighting), there is some tendency for these worries to color our imaginative representation of the thinking of those more naïve subjects. But with these natural distortions of the imagination understood, we can see a single standard as measuring the naïve agent who considers minimal evidence and the anxious agent who reaches for more.²

² Thanks to Diana Raffman, Patrick Rysiew, Sergio Tenenbaum and Jonathan Weisberg for feedback on an earlier draft.

References

- DeRose, K. (2009). *The Case for Contextualism: Knowledge, Skepticism, and Context, Volume 1*. New York: Oxford University Press.
- Kruglanski, A. W., & Webster, D. M. (1996). Motivated closing of the mind: "Seizing" and "Freezing". *Psychological review*, *103*(2), 263-283.
- Lerner, J. S., & Tetlock, P. E. (1999). Accounting for the Effects of Accountability. *Psychological Bulletin*, *125*(2), 255-275.
- Payne, J., Bettman, J., & Johnson, E. (1993). *The Adaptive Decision Maker*. New York: Cambridge University Press.
- Weatherson, B. (2005). Can we do without Pragmatic Encroachment? *Philosophical Perspectives*, *19*(1), 417-443.
- Williamson, T. (2005). Contextualism, Subject-Sensitive Invariantism and Knowledge of Knowledge. *The Philosophical Quarterly*, *55*, 213-235.