

To appear in *Routledge Handbook of Epistemic Contextualism*, J. Ichikawa, ed.
Jennifer Nagel and Julia Jael Smith
April 22, 2016

The Psychological Context of Contextualism

1. Introduction

Philosophy is not the only discipline concerned with shifting epistemic intuitions. Psychology also studies intuitive impressions of knowledge, often with an eye to the ways in which these intuitive impressions can shift as circumstances change. Cognitive and social psychologists explore factors figuring prominently in the contrasting cases motivating contextualism, including rising and falling practical stakes, heightened self-consciousness, and worries about the possibility of error. In some cases, these factors are taken to distort perceptions of knowledge in the attributor, causing the impression of a shift in knowledge where no real shift in knowledge has occurred. In other cases, these factors are taken to have an impact on the subject's actual possession of knowledge, for example by raising or lowering the subject's confidence. Where attributor perceptions of knowledge march in step with variations in the subject's knowledge, there is no reason to consider a shift in perceptions mistaken, but such shifts may risk being misunderstood by theorists who are unaware of the subtle ways in which contextual circumstances change a subject's cognition. If the key claim of contextualism is that statements of the form "S knows that p" can express propositions differing in truth value for subjects matched on traditional epistemic factors (such as confidence and accuracy), philosophers need to be aware of the subtle ways in which these traditional epistemic factors can change, and be instinctively registered as changing, in response to changing circumstances. Equally, philosophers who are sympathetic to contextualism will want to know how it might be psychologically possible for us to deploy shifting standards in evaluating others.

This article reviews psychological work relevant to the attainment and perception of knowledge for subjects and attributors in shifting circumstances, and discusses the upshot of this work for epistemic contextualism. One major body of relevant research is dual process theory, according to which there are two

broad types of judgment, intuitive and reflective. Section 2 outlines the contrast between these modes of thought, highlighting features of the contrast that matter to the production and perception of knowledge in changing circumstances. Beyond this basic binary division, there are more finely graded divisions between more and less effortful ways of thinking. Because the cases motivating contextualism often contrast subjects who seem to need more or less information in order to make up their minds, epistemologists can benefit from reviewing the recent literature on variations in information search, as we do in Section 3. People make up their minds in various ways, sometimes deciding on the basis of a glance or a single evidential cue, sometimes after extended evidence collection. This section looks at the factors governing search termination, and in particular, at factors affecting the relationship between subjects' confidence and their information search behavior. The perspective of the attributor of knowledge is examined more closely in Section 4. Ordinary people track states of knowledge and belief intuitively, but imperfectly. This section gives a quick overview of empirical work on mindreading, or the capacity to attribute mental states, and then takes a closer look at the most-studied imperfection in mindreading, the bias of epistemic egocentrism. Epistemic egocentrism is a natural limitation in our evaluation of more naïve subjects. This robust bias in perspective-taking inclines us to project certain features of our own predicament onto others, even when we are well aware on reflection that this is inappropriate. It is possible that the bias of epistemic egocentrism explains some of the asymmetries in patterns of intuition motivating contextualism; Section 5 will examine this possibility, while also reviewing other ways in which epistemologists have tried to use psychological research to explain, and in some cases explain away, the intuitions motivating contextualism. A brief concluding section summarizes our findings and discusses questions for future research.

2. Dual process theory

Over the past forty years, a number of originally independent research programs in several branches of psychology have converged on the idea that there is a significant divide between two core ways of making judgments. Following Jonathan Evans and Keith Stanovich (2013), we will label these intuitive (or Type 1) and

reflective (or Type 2) thinking. A cluster of characteristics is associated with each side of the contrast: intuition is said to be unconscious, low-effort, rapid and holistic, while reflection is conscious, effortful, slow and analytic (Evans 2008). It is now increasingly accepted that the key difference between these two modes of thought is whether key processing occurs in working memory, which is always consciously accessible (Evans and Stanovich 2013, Carruthers 2015). To take examples commonly used to illustrate this contrast, the intuitive process of face recognition happens on the basis of geometrical calculations that are not consciously available; the reflective process of mental multiplication of two-digit numbers happens on the result of conscious sequential digit-by-digit operations. Because working memory has a limited capacity, when it is burdened with another task (such as reciting the alphabet backwards), intuitive tasks are relatively unimpaired but reflective tasks become very difficult.

The reflective mode of thought is engaged in hypothetical thinking, high-stakes thinking, and almost all exercises of negation (Evans, 2008; Strack & Deutsch, 2004). So for example, in order to evaluate the truth value of the proposition “it is not the case that the Pope has just died”, a sequence of thoughts is necessary: one needs to entertain the embedded positive proposition about the Pope’s death, and then negate it. A question about the name of the current Pope, by contrast, would for someone with the relevant background knowledge simply bring an answer to mind without consciously accessible sequential thought. Some propositions involve operations which demand reflective cognition: “it is not the case that this is a cleverly disguised mule”, for example, can only be judged reflectively. Other propositions (“this is a zebra”) may be judged either intuitively or reflectively, depending on one’s circumstances and motivations.

Ordinarily, shifts between intuitive and reflective thought are automatic and silent: we do not typically decide to think reflectively, or focus on which way we are thinking at any given time. Once a person has entered the reflective mode of thought and consciously entertained a particular content (for example, a scenario involving a painted mule), it will however take some time for this content to fade from working memory, and it will remain active as an input to one’s thinking on any related question (Carruthers, 2015). As section 4 of this chapter will explore in more detail, related questions could include not only one’s own judgments about

the animal but also one's representation and evaluation of the judgments of others. The person who is now thinking reflectively about possible painted mules will go back to making naïve intuitive judgments about zebras (and about others' judgments about zebras) only when that possibility fades from working memory, and not at will.

The line between knowledge and ignorance seems to run orthogonal to the line between intuitive and reflective thought. Both automatic exercises of face recognition and effortful exercises of mental arithmetic are ordinarily seen as issuing in knowledge, but either type of process can also occasionally lead to error or mere belief. It's not obvious that intuitive judgments are generally made to a lower standard than reflective ones, or that they have any lesser claim to knowledge: indeed, given that reflective judgments are typically composed of sequences of intuitive judgments (for example, a sequence of intuitive digit-by-digit manipulations in reflective arithmetical thinking), the epistemic quality of reflective judgments presumably depends on the quality of underlying intuitive judgments. A switch from intuitive to reflective judgment may however bring with it a demand for different evidential resources to settle a question; the consequences of this point for contextualism will be examined in section five.

3. Search Termination

The binary contrast between intuitive and reflective thought is not the only type of variability in judgment. Within the domain of reflective thought, we may engage in more or less extensive searching and weighing of evidence. In theory, we might continue searching and weighing evidence indefinitely on any given question, never reaching an answer. In practice, we need to stop at some point, typically without conscious deliberation about where that point will be. Indeed, if we always needed to preselect a stopping point before making up our minds on any given question, we'd be trapped in a regress of thinking about how hard we needed to think. Cognitive effort is, fortunately, governed by an array of instinctive mechanisms. Unsurprisingly, these mechanisms are broadly adaptive, reflecting the costs and benefits of evidence acquisition and accuracy (Payne, Bettman, & Johnson, 1993). We also seek more evidence when we are made to feel self-conscious about our

choices, and when we anticipate needing to justify them to others (Lerner & Tetlock, 1999). Changes in perceived stakes have a natural impact on confidence, so high-stakes subjects need more evidence to make up their minds (Hausmann & Lage, 2008). There are many theories about the relationship between confidence and knowledge, but it is fairly uncontroversial that at a minimum, a subject needs to have made up her mind about p in order to know that p , and there is an objective psychological fact of the matter as to whether the subject has made up her mind on any given point. Some theorists—notably Keith DeRose (2009)—hold that the possession of knowledge may require higher levels of confidence than mere belief. Subjects with perceived high stakes also adopt (and are also instinctively expected to adopt) more elaborate and effortful decision-making strategies (McAllister, Mitchell, & Beach, 1979). Given that one’s confidence and one’s way of weighing evidence are traditional factors in knowledge, the contrasting high-stakes and low-stakes subjects in the scenarios taken to motivate contextualism are not typically matched on traditional factors, even when it is specified that they have the same evidence set in mind when making a judgment.

There are several ways of measuring the relationship between stakes and evidence collection. One influential research program describes evidence collection as driven by “variable confidence thresholds”: depending on our environment, we will require a higher or lower degree of confirmation from our evidence in order to make up our minds: where decision problems have proved difficult we come to require higher confidence in order to settle a question. According to this program, the level at which these confidence thresholds are set is something we ordinarily learn in decision environments, alongside our learning about environmental features. In a climate where accurate decisions can be made with fairly minimal search, we learn to stop searching for information fairly rapidly; where more exhaustive search is required for accuracy, we start to need much more confidence to make up our minds and call off the hunt for evidence (Lee, Newell, & Vandekerckhove, 2014). Empirical research into threshold-setting is carried out in controlled laboratory conditions where participants make hundreds of consecutive judgments, often with feedback on accuracy, and with evidence artificially manipulated to require various levels of evidence search to achieve accuracy in various batches of these judgments. It is an interesting and open question what works in practice as a “decision environment” for setting a confidence threshold, for the purpose of intuitive evaluation of epistemological

scenarios. Taking as a baseline a friendly decision environment (a normal zoo, where the task is animal identification), merely imagining a hostile decision environment (a zoo run by practical jokers), may be sufficient to trigger a higher threshold and the craving for additional evidence. In any event, shifting confidence thresholds are part of the invisible psychological baggage of judgment: it is easy to see the overt cues on which judgments are made, and these are typically spelled out in the contrasting scenarios we are invited to evaluate, but we are also moved by inner factors controlling our sensitivity to those cues. Because confidence is a traditional factor determining whether someone should count as knowing, when epistemologists are deciding whether contrasting characters in epistemological scenarios are well-matched on traditional factors, variable confidence thresholds should be taken into account.

4. Mindreading and egocentric bias

Intuitions about the presence and absence of knowledge feature prominently in our everyday social navigation: verbs meaning ‘know’ appear in all natural languages, and the capacity to recognize knowledge in others develops early and is used often (Bartsch & Wellman, 1995; Goddard, 2010). Commonly labelled our ‘mindreading’ capacity, the intuitive ability to attribute states of knowledge, belief and desire is active not only in live encounters with other agents, but also in watching videos and reading scenarios about them: indeed, mindreading works the same way—and shows the same limitations—across these various platforms (Saxe, 2006). There is nothing peculiarly philosophical about having an impression that an agent in a story has or lacks knowledge; indeed, the pattern of intuitions taken to motivate contextualism can also be detected in undergraduates without philosophical training (Nagel, Juan, & Mar, 2013). Mentioning a possibility of error (“sometimes clocks are broken”) very significantly tips an audience towards using “thinks” rather than “knows” in describing an agent’s state of mind (for example, towards describing someone who glances at a clock as having just a belief, rather than knowledge, about the current time, even when it is stipulated that the clock is accurate).

Mindreading is an intuitive capacity in the sense that it performs calculations that are not consciously available to the mindreader: we watch another person move her hand towards an object, and we instinctively see her as knowing that the object is there, and as wanting it. In attributing those mental states we rely on some subtle combination of perceptible cues in the person's posture and orientation towards the scene; likewise, when we read a story about a person who is making a decision or statement, we register various features of their described predicament, actions and assertions as relevant to whether they possess knowledge or mere belief. To the extent that the line between knowledge and belief is drawn intuitively, even in highly novel cases, we have limited introspective access to the rules we follow in drawing this line. If it were introspectively obvious why we saw any given case as involving knowledge, we would not have such lively controversies over the reasons why we see shifting attributions in response to various pairs of cases—for example, over whether we are really responding to stakes, or to the salience of error possibilities. To discover what aspects of a scenario matter to knowledge attribution, it seems we have to experiment with a great variety of cases in which various factors are manipulated systematically. We also have to be alert to the limits of our powers of stipulation: we need to be cautious, for example, about taking the traditional factor of confidence to be set as equal across a pair of contrasting cases just by a stipulation that the subject in the second case is as confident as the subject in the first. If our intuitive attributions of confidence are naturally driven by behavioral cues such as whether the individual is seeking further evidence on a question, then we may instinctively register the searching agent as less confident than the settled agent, even if it has been stipulated that they are equally confident.

A further difficulty in relying on intuitive evidence is that intuitive mechanisms have various natural limitations: given the boundless variety of ways in which states like knowledge can be manifested, our mechanisms for the intuitive detection of knowledge need to make certain assumptions and take certain shortcuts. Just as our visual mechanisms can be tricked by various patterns of cues to depth, magnitude and motion, producing visual illusions, so also our mindreading mechanisms are subject to structurally similar limits, producing cognitive illusions. The deepest bias in mental state attribution is to assume one's own frame of mind as a default for others: known most generally as “epistemic egocentrism”, this bias affects both children and adults (Birch & Bloom, 2004; Royzman, Cassidy, & Baron, 2003). There is an asymmetry in this bias: it impairs

our evaluation of others who are seen as more naïve, but not our evaluation of others who are seen as better-positioned than we are. The bias persists even in the face of forewarnings and incentives to cancel it, and even when we are explicitly aware that the agent we are evaluating does not share our privileged knowledge or concerns. The existence of this bias should raise concerns about the accuracy of our responses to ignorant high stakes cases, for example, where it is stipulated that the naïve agent is in a life-or-death situation but unaware of this fact. Egocentric bias would incline us to expect such an agent to take more care or collect more evidence than his low-stakes counterpart, before making up his mind.

Epistemic intuitions are ordinarily generated to help us interact with other agents, making sense of their actions and assertions. It makes sense that epistemic intuitions can be elicited only by imagining the target agent as acting or asserting something: the question “does Smith know where his car is?” invites contemplation of a scenario in which Smith needs to drive somewhere, or a scenario in which he is asked about his car. Even if knowledge is a standing state of a subject that persists through sleep and inactivity, if we try to restrict ourselves to imagining Smith as silent and inactive, perhaps deeply asleep, it is hard to generate an intuitive answer to the question of whether he now knows where his car is. The need to imagine a concrete situation may raise problems for our readings of contrasting cases in which we evaluate what speakers say about an absent third party whose strength of epistemic position is supposed to be held fixed across the contrast. Keith DeRose’s (2009) Lena cases may provide example of this kind of difficulty. The cases concern an office worker (Lena), who has a modest amount of evidence that her co-worker (John) was at the office today: she has seen his hat on a hook in the hallway, and she has overheard a conversation suggesting he was present. Lena’s epistemic situation is later evaluated by a colleague (Thelma) in a low-stakes barroom conversation about John’s whereabouts, and by another colleague (Louise) in an interrogation with the police about a serious crime. When we think about Louise’s predicament, as she is being questioned by the police about what Lena knows, she would naturally anticipate that a positive answer—“yes, Lena knows whether John was there”—would lead the police to interrogate Lena. In her low stakes conversation, Thelma wouldn’t need to imagine Lena under interrogation: she could instead ask herself about what Lena would be likely to say if she were also present in the casual barroom. If these contrasting cases naturally trigger us to imagine different types of scenario across

the contrast—Lena being questioned by police in one case, Lena engaged in casual conversation in the other—then our evaluation of the knowledge attributions will be a function of our assessment of Lena’s resources to make up her mind in those different scenarios. If we stipulate that in both cases, Lena is not in fact considering the question of John’s whereabouts at the moment, this does not ensure that we will imagine her walking along with her mind on other matters as we try to assess the propriety of various claims about what she knows. We may need to imagine her as acting or asserting something relevant to a proposition in order to get an intuitive sense of whether she knows it; if so, then we need to be careful about what attributors would naturally imagine her as doing or saying, when evaluating her from different perspectives.

5. Applications of psychological research to epistemic contextualism

Philosophers debating the merits of contextualism have engaged with psychological work in a variety of ways. To illustrate their differences, it will be useful to consider a case, drawn from Jonathan Vogel (1990). Smith accurately remembers having parked his car in the lot an hour ago; does Smith know that his car is parked in the lot? It seems natural to say that he does know. But when we make salient an alternative possibility—that many cars are stolen each day, and that Smith’s car could have been stolen in the past hour—it seems natural to judge that Smith doesn’t know that his car is parked in the lot. In this case nothing about Smith’s epistemic position changes when the alternative possibility is introduced in the voice of the narrator, and yet mentioning the possibility of error affects the reader’s intuition about whether he possesses knowledge. In general, we tend to think that a subject has less of a claim to knowledge when possible alternatives to the key proposition are mentioned.

Timothy Williamson and John Hawthorne have both argued that a form of bias known as the availability heuristic interferes with our ability to make an accurate judgment in cases in which alternative possibilities are mentioned (Hawthorne, 2004a, 2004b; Williamson, 2005). The availability heuristic is a phenomenon whereby we overestimate the likelihood of events that can be easily imagined and recalled (Tversky & Kahneman, 1973). When we are judging whether a subject has knowledge and alternative

possibilities are mentioned, we overestimate the likelihood of these possibilities, which causes us to mistakenly judge that the subject lacks knowledge. However, the availability heuristic explanation is problematic for a number of reasons. The availability heuristic leads us to judge probabilities on the basis of ease (or difficulty) of recall or imagination, but it does not always inflate these probabilities: discussion of easy-to-imagine scenarios makes them seem more likely, but discussion of hard-to-imagine scenarios makes them seem less likely (Sherman, 1985). It is not clear that the problematic scenarios mentioned by epistemologists invariably count as easy to imagine, in the relevant sense. It is also not clear that a single explicit mention of any alternative possibility is enough to activate the availability heuristic; indeed, some research suggests that while extended exposure to a type of scenario (such as repeated media reports of airplane crashes) will increase one's estimation of its likelihood, but a single mention of a scenario (for example, a novel situation involving tricky lighting) would have the opposite effect (Oppenheimer, 2004). Meanwhile, reliance on the heuristic is attenuated or cancelled if we are forewarned of it (Schwarz, 1998) or in conditions where we expect to have to defend our judgments to others (Lerner & Tetlock, 1999); it is unclear how so fragile a bias could continue to produce the intuitions driving skepticism and contextualism for forewarned participants in an epistemology seminar room. Lastly, even if it were the case that availability bias operated as Williamson and Hawthorne propose, a more detailed explanation is still needed for why overestimating the probability of alternative possibilities would make us more hesitant to ascribe knowledge (Nagel, 2010).

The bias of epistemic egocentrism has been offered as another explanation of the shift. Applied to the car case, if we are egocentrically evaluating Smith as if what is salient to us is also salient to him, we will mistakenly evaluate him as though he were also concerned about the possibility of auto theft. If we are imagining Smith as being actively worried about the possibility of theft, we will expect him to secure some evidence that tells against that possibility before arriving at the judgment that his car is parked in the lot after all. When Smith fails to collect this evidence, we see his judgment that his car is parked in the lot as being made too hastily or carelessly to count as knowledge. On this picture, provided that Smith can in fact know that his car is in the lot, our evaluation that Smith lacks knowledge is in error, since it is based on a representation of Smith's mental states that is the natural product of a mental state attribution bias, a representation that would not in fact fit the

real mental states of a person who occupied the position Smith is described as occupying in the scenario (Nagel, 2010).

If the egocentrism hypothesis is correct, we should find that if the subject is described as being aware of the possibility of error may not greatly alter how willing we are to attribute knowledge to him. If egocentrism already causes us to misjudge Smith as being concerned with auto theft, adding to the story that Smith is aware of the possibility of auto theft shouldn't change our assessment of what he knows. If the egocentrism hypothesis is correct, we should also find that if the subject is described as entertaining the possibility that he is wrong, this makes little difference to how willing we are to attribute knowledge to him. If egocentrism already causes us to misjudge Smith as being concerned with auto theft, adding to the story that Smith is in fact entertaining the possibility that his car may have been stolen should again not affect our assessment of what he knows. Alexander, Gonnerman, and Waterman test these two predictions, finding, as expected, no significant difference between participants' willingness to attribute knowledge in the relevant types of case (Alexander, Gonnerman, & Waterman, 2014). In addition, they observe a modest correlation between the degree to which participants saw subjects as sharing their concerns, and the degree to which participants denied that the subject possessed knowledge – again confirming that the effects of epistemic egocentrism are at play in accounting for our tendency to judge that subjects fail to possess knowledge in the relevant cases.

These findings offer some additional support for the egocentrism hypothesis. However, some lingering concerns with this line of explanation remain. First, in Alexander, Gonnerman, and Waterman's results, there is only a modest correlation between the degree to which subjects were seen as sharing the attributor's concerns and knowledge denials. This suggests that egocentric bias is only a part of the explanation for why we are inclined to deny knowledge when alternative possibilities are made salient, and that there are other factors at play that contribute to this particular pattern of intuitions. More problematic is the fact that, if the egocentrism hypothesis is correct, then we should judge subjects in the relevant cases not only as failing to know, but also as failing to be justified in their beliefs. If we judge Smith's belief as falling short of knowledge because we (mistakenly) believe he has formed it hastily or carelessly, it seems to follow that we should also see his belief as being unjustified for the very same reason. However, this prediction is not borne out: in cases where the

possibility of error is mentioned, people typically perceive a subject as possessing justified true belief but not knowledge (Nagel et al., 2013).

A third psychological explanation involving variations in cognitive effort may provide a better way forward. Simple and familiar problems are typically handled intuitively, where complex problems—such as those including new situations, hypothetical scenarios, or involving a negated phrase—demand reflective thought, and more extensive evidence search. The routine judgment that Smith makes when he utters “I know my car is parked in the lot” on the basis of his memory of having parked it there twenty minutes ago would be arrived at heuristically, using a ‘low’ strategy. Before the possibility of error is mentioned, an attributor who is seeing Smith’s circumstances accurately will also most naturally judge that Smith knows his car is parked in the lot through a low cognitive strategy.

Descriptions of agents’ situations that mention the possibility of error, on the other hand, characteristically contain the hallmark features that precipitate a higher cognitive strategy: the presence of a hypothetical alternative; an unfamiliar problem; often too the presence of negation. So it is quite plausible that the mention of the possibility of car theft causes the attributor to shift into a more effortful mode of thought—more effortful than the one used by the subject himself. We may indeed be saddled with this mode of thought as long as the problematic theft scenario remains active in working memory for us. When we are actively conscious of the possibility of auto theft, it is natural to feel that if we were in Smith’s position, we could not know that the car is parked in the lot without collecting evidence that would rule out the possibility that it has been stolen. Now, when we are asked to evaluate whether *Smith* knows that his car is parked in the lot, we conclude that Smith doesn’t know, because we use our higher threshold as a benchmark for evaluating Smith’s reasoning (Nagel, 2011). Unlike the egocentrism hypothesis, the dual process explanation does not require that we think of Smith as actually engaging in more effortful thinking. It merely requires that we sanction Smith for failing to collect more evidence, or for failing to think harder about the situation, without seeing him as actually sharing our concerns. On this picture, we judge that Smith lacks knowledge because we believe that he has not collected the quantity of evidence that is necessary for knowledge in his circumstances; but importantly, we are still able to judge Smith’s belief as justified, since we can still represent him as forming his belief through a low

cognitive strategy which is perfectly natural for him given his circumstances. When we think about whether his failure to collect additional evidence might affect the justification of his belief, we can recognize that he had no special reason to think that additional evidence collection could be appropriate in this situation.

Drawing on many of the same resources from the psychological literature, Mikkel Gerken provides an alternate explanation for our tendency to deny knowledge when alternative possibilities are mentioned (Gerken & Beebe, 2014; Gerken, 2013). Gerken's *epistemic focal bias* account relies on two underlying principles: (i) normally, an agent processes a salient alternative possibility to a subject's knowledge that *p* as epistemically relevant; and (ii) normally, attributors form judgments about a subject's claim to knowledge on the basis of a *prima facie* reason that constitutes a limited amount of the total evidence (Gerken, 2013). The first principle is supported by the general phenomenon of *focal bias*: typically, we tend to give judgments that place undue weight on what occupies our focus, because we typically process what is salient to us even when what happens to be salient is irrelevant to the task at hand (Gigerenzer & Todd, 2000; Stanovich, 2009). The second 'satisficing' principle is an epistemological specification of a more general observation prevalent in the psychological literature (Jonathan Evans, 2008; H. A. Simon, 1955; H. Simon, 1990; Stanovich, 2009): namely, that we conserve cognitive effort in arriving at judgments by settling them on a basis of a *prima facie* reason before we have processed all the evidence. Given these two principles, it is possible to explain the shift in our intuitions about a subject's knowledge that occurs when a possible alternative is mentioned as a product of a bias that occurs when a contextually salient alternative possibility is in fact epistemically irrelevant.

On this account, when we are told that Smith's car could have been stolen in the past twenty minutes, we regard this information as epistemically relevant due to (i). Moreover, that Smith's car could have been stolen in the past twenty minutes without Smith's knowing it counts as a *prima facie* reason to think that Smith doesn't know that his car is parked in the lot. Given this *prima facie* reason, in accordance with (ii), we judge that Smith fails to possess knowledge. However, this conclusion is incorrect if the information that Smith's car could have been stolen is not epistemically relevant—as certain invariantists would maintain.

Unlike contextualists, advocates of the variable cognitive effort explanation and the epistemic focal bias explanation take the introduction of a salient alternative to be a red herring that is not in fact a relevant

consideration when it comes to whether the subject is rightly said to know. On the variable cognitive effort model, introducing the salient alternative can change the attributor's frame of mind in a way that will make it harder for her to see the naïve subject as knowing: she is right that if Smith started to think about the possibility that his car had just been stolen, he would have to collect more evidence before making a judgment about its current location, but mistaken to evaluate Smith as if he shared her higher standard for evidence collection. On the epistemic focal bias model, the salient alternative is simply a distraction: neither Smith nor the attributor should contemplate the possibility of recent theft, in this case, and it would be a mistake to do so, either in thinking about the location of the car or in thinking about what Smith knows.

Both hypotheses provide psychological explanations that allow the strict invariantist to explain the patterns of intuition that have often been taken to support contextualism. But both hypotheses also require further epistemological commitments to complete the invariantist's argument. For example, the proponent of the epistemic focal bias account must supply a convincing story about why contextually relevant alternatives are not always epistemically relevant, and further details about what constitutes an epistemically relevant alternative. The invariantist proponent of the dual process explanation must provide a more complete story about when judgments formed through more and less effortful cognitive strategies constitute knowledge.

Although psychological explanations provide a promising framework for strict invariantists to explain the linguistic data that is often taken to support contextualism, these explanations could also admit of interpretations that are consistent with a contextualist's theoretical commitments. Consider the dual-process explanation, for example. The contextualist might claim that cases in which we use higher cognitive strategies correspond to higher-standards contexts for knowledge attributions, while cases in which we use lower cognitive strategies correspond to lower-standards contexts for knowledge attributions, and that both kinds of contexts allow for true attributions of knowledge.

Nevertheless, psychological explanations of our intuitions about knowledge ascriptions in cases involving the possibility of error arguably undermine contextualist's arguments for the context-sensitivity of "knows" that rely on these patterns of intuition. For if psychological explanations of these patterns reveal that

the intuitive data are *prima facie* equally compatible with both invariantism and contextualism, the burden of proof is on the contextualist to show why the context-sensitivity of “knows” is still the best explanation for this data, given that positing context-sensitivity here also demands a more elaborate semantics.

6. Conclusion and questions for future research

Thought experiments involve an interplay between what is explicitly stated in the scenario and the resources we apply to them. There would be no interest in asking a question at the end of a scenario which simply prompted one to repeat back one of the scenario’s explicit stipulations; we enjoy a sense of discovery in our responses partly because we cannot always tell in advance how our intuitive mechanisms will respond. Epistemological scenarios don’t need to stipulate overtly whether someone judges intuitively or reflectively, or what their evidence search termination points or confidence threshold would be; we have intuitive functions that do this work for us, automatically and silently. It is easy to overlook the significance of these silent factors when we are diagnosing the significance of our intuitive reactions to particular cases, but important to take them into account when we are figuring out what these cases are telling us about our knowledge-ascribing practices, and (perhaps indirectly) about knowledge itself.

According to contextualists, our use of “knows” is governed by an appreciation of the target’s “strength of epistemic position” (to use Keith DeRose’s expression for the traditional factors of evidence, confidence and the like), as measured against whatever epistemic standard is operative in the relevant conversational context. Our grasp of both of these notions can be improved by studying natural variations in judgment and evidence collection on the one hand, and natural variations and limitations in epistemic assessment, on the other. To the extent that contextualism is motivated by patterns in ordinary language, exploring the

psychological underpinnings of these patterns can help us see whether contextualism is the right way to explain them.¹

References:

- Alexander, J., Gonnerman, C., & Waterman, J. (2014). Salience and Epistemic Egocentrism: An Empirical Study. In J. R. Beebe (Ed.), *Advances in Experimental Epistemology* (pp. 97–118). New York: Bloomsbury.
- Bartsch, K., & Wellman, H. M. (1995). *Children talk about the mind*. New York: Oxford University Press.
- Birch, S., & Bloom, P. (2004). Understanding children's and adults' limitations in mental state reasoning. *Trends in Cognitive Sciences*, 8(6), 255–260.
- Carruthers, P. (2015). *The centered mind: what the science of working memory shows us about the nature of human thought*. Oxford University Press.
- DeRose, K. (2009). *The Case for Contextualism: Knowledge, Skepticism, and Context, Volume 1*. New York: Oxford University Press.
- Evans, J. (2008). Dual-processing accounts of reasoning, judgment, and social cognition. *Annual Review of Psychology*, 59, 255–278.
- Evans, J., & Stanovich, K. (2013). Dual-Process Theories of Higher Cognition Advancing the Debate. *Perspectives on Psychological Science*, 8(3), 223–241.
- Gerken, M. (2013). Epistemic Focal Bias. *Australasian Journal of Philosophy*, 91(1), 41–61.
- Gerken, M., & Beebe, J. R. (2014). Knowledge in and Out of Contrast. *Noûs*, 49(2), n/a–n/a.
- Gigerenzer, G., & Todd, P. M. (2000). *Simple Heuristics that Make Us Smart*. New York: Oxford University Press.
- Goddard, C. (2010). Universals and variation in the lexicon of mental state concepts. In B. Malt & P. Wolff (Eds.), *Words and the Mind* (pp. 72–93). New York: Oxford University Press.
- Hausmann, D., & Lage, D. (2008). Sequential evidence accumulation in decision making: The individual desired level of confidence can explain the extent of information acquisition. *Judgment and Decision Making*, 3(3), 229–243.
- Hawthorne, J. (2004a). *Knowledge and Lotteries*. New York: Oxford University Press.
- Hawthorne, J. (2004b). Replies. *Philosophical Issues*, 14, 510–23.
- Lee, M. D., Newell, B. R., & Vandekerckhove, J. (2014). Modeling the adaptation of search termination in human decision making. *Decision*, 1(4), 223–251.

¹ For helpful comments on an earlier draft of this chapter, we would like to thank Nathan Cockram, Jonathan Ichikawa, and Micah Smith. We would also like to acknowledge the support of the Social Sciences and Humanities Research Council of Canada.

- Lerner, J. S., & Tetlock, P. E. (1999). Accounting for the Effects of Accountability. *Psychological Bulletin*, 125(2), 255–275.
- McAllister, D. W., Mitchell, T. R., & Beach, L. R. (1979). Contingency-Model for the Selection of Decision Strategies - Empirical-Test of the Effects of Significance, Accountability, and Reversibility. *Organizational Behavior and Human Performance*, 24(2), 228–244.
- Nagel, J. (2010). Knowledge Ascriptions and the Psychological Consequences of Thinking About Error. *Philosophical Quarterly*, 60(239), 286–306.
- Nagel, J. (2011). The psychological basis of the Harman-Vogel paradox. *Philosophers' Imprint*, 11(5), 1–28.
- Nagel, J., Juan, V. S., & Mar, R. A. (2013). Lay denial of knowledge for justified true beliefs. *Cognition*, 129(3), 652–661.
- Oppenheimer, D. (2004). Spontaneous discounting of availability in frequency judgment tasks. *Psychological Science*. Retrieved from <http://pss.sagepub.com/content/15/2/100.short>
- Payne, J. W., Bettman, J. R., & Johnson, E. J. (1993). *The Adaptive Decision Maker*. New York: Cambridge University Press.
- Royzman, E. B., Cassidy, K. W., & Baron, J. (2003). “I know, you know”: Epistemic egocentrism in children and adults. *Review of General Psychology*, 7(1), 38.
- Saxe, R. (2006). Why and how to study Theory of Mind with fMRI. *Brain Research*, 1079(1), 57–65.
- Schwarz, N. (1998). Accessible content and accessibility experiences: The interplay of declarative and experiential information in judgment. *Personality and Social Psychology Review*, 2(2), 87.
- Sherman, S. (1985). Imagining can heighten or lower the perceived likelihood of contracting a disease: The mediating effect of ease of imagery. *Personality Social Psychology Bulletin*.
- Simon, H. (1990). *Reason in Human Affairs* (1 edition). Stanford, Calif: Stanford University Press.
- Simon, H. A. (1955). A Behavioral Model of Rational Choice. *The Quarterly Journal of Economics*, 69(1), 99–118. <http://doi.org/10.2307/1884852>
- Stanovich, K. (2009). Distinguishing the reflective, Algorithmic, and Autonomous Minds: Is It Time for a Tri-Process Theory? In J. Evans & K. Frankish (Eds.), *In Two Minds: Dual Process and Beyond*. Oxford: Oxford University Press.
- Strack, F., & Deutsch, R. (2004). Reflective and impulsive determinants of social behavior. *Personality and Social Psychology Review*, 8(3), 220–247.
- Tversky, A., & Kahneman, D. (1973). Availability: A heuristic for judging frequency and probability. *Cognitive Psychology*, 5(2), 207–232. [http://doi.org/10.1016/0010-0285\(73\)90033-9](http://doi.org/10.1016/0010-0285(73)90033-9)
- Vogel, J. (1990). Are there counterexamples to the closure principle? In M. Roth & G. Ross (Eds.), *Doubting: Contemporary Perspectives on Skepticism* (pp. 13–27). Dordrecht: Kluwer.
- Williamson, T. (2005). Contextualism, Subject-Sensitive Invariantism and Knowledge of Knowledge. *Philosophical Quarterly*, 55(219), 213–235.