



# Epistemic Intuitions

Jennifer Nagel\*

University of Toronto

---

## Abstract

We naturally evaluate the beliefs of others, sometimes by deliberate calculation, and sometimes in a more immediate fashion. Epistemic intuitions are immediate assessments arising when someone's condition appears to fall on one side or the other of some significant divide in epistemology. After giving a rough sketch of several major features of epistemic intuitions, this article reviews the history of the current philosophical debate about them and describes the major positions in that debate. Linguists and psychologists also study epistemic assessments; the last section of the paper discusses some of their research and its potential relevance to epistemology.

---

When Plato's character Theaetetus proposes defining knowledge as true belief, Socrates tells him a story about a talented lawyer who had to defend a difficult case. Facing charges arising from some violent incident, the lawyer's client could provide him with little firm evidence concerning the event, and no neutral eyewitnesses to call. The client actually happened to be innocent. However, since the lawyer had limited time to present his case in court, he decided that charming the jury with rhetoric would probably be more effective than trying to instruct them about what really happened. This tactic worked well, and the jury became convinced that the defendant was innocent. Responding to Socrates' story, Theaetetus readily grants that in this case the members of the jury do not *know* that the defendant is innocent, despite having a true belief on this point. Socrates and Theaetetus both take this case to refute the theory that knowledge is simply true belief (Plato 201a–c).

If you also find it easy and natural to evaluate the jury in Plato's story as lacking knowledge, you can be described as sharing Theaetetus's epistemic intuition on this case. You might wonder what enables you to respond to the story as you do, and whether this response tells you something about knowledge itself, or whether it is better seen as telling you something about your existing concept or tacit theory of knowledge. You might wonder about the range of cases that can spark such reactions, and whether these reactions would have to be the same for all people, or even the same for you in different contexts. You might also wonder about

the degree to which it is possible, or desirable, for a theory of knowledge to accommodate such evaluations, and whether some such evaluations should be given more weight than others. Questions of this type have gained new prominence as analytic epistemology has become more self-conscious about its methods, and contemporary epistemologists are presently far from consensus on how these questions should be answered. Section one of this article outlines several fairly uncontroversial features of epistemic intuitions, and discusses several accounts of the relationship between intuitions and theories. Section two explains the historical background behind the current debate about intuitions, and section three maps out the major positions taken by philosophers in this debate. Intuitions about knowledge and justification are also studied by linguists and psychologists; section four discusses some empirical findings about epistemic intuitions, and suggests some ways philosophers might make use of these results.

### *1. Introduction: A Rough Sketch of Epistemic Intuition*

Theaetetus' evaluation of the jury's belief is a prime example of the sort of epistemic intuition that will interest us in what follows. The expression 'epistemic intuition' is sometimes used very broadly, as a label for any immediate (or not explicitly inferential) assessment of any claim of interest to epistemologists; for example, philosophers sometimes speak of having an intuition that knowledge entails belief, or that knowing something entails that one couldn't easily have been wrong about it. The nature and basis of these more general judgments is an interesting (and large) topic in its own right, as is the even broader question of the value of intuitions in philosophical debates outside of epistemology. The present article restricts its focus to epistemic evaluations of particular cases, although attention will be paid to the ways in which particular case evaluations interact with more general principles. Evaluations of the particular kind that will matter here occur when one is presented with a description of some real or hypothetical subject's situation, where the description doesn't already label the subject's condition as one of knowledge, or justified belief, or whatever epistemological notion is being tested. Clear epistemic intuitions arise when the subject's described condition plainly appears to fall on one side or the other of some significant divide in epistemology, such as the divide between knowledge and ignorance, or justified and unjustified belief.

The word 'appears' in that last sentence deserves emphasis: the metaphor of appearance is used heavily in discussions of epistemic intuition, and several distinct facets of this metaphor should be separated. First, the way that epistemic intuitions strike us has something in common with the way that visual appearances strike us: at least in clear cases like the one Plato gives us, even if effort is required to understand the story,

production of the epistemic intuition itself does not seem to require effort, and does not feel like the result of a deliberate cognitive process like guessing or calculation. Rather, the judgment that the jury lacks knowledge has a certain immediacy, like a simple perceptual judgment. In particular, given the way it runs against his explicit working theory of knowledge, Theaetetus' intuition that the jury lacks knowledge cannot be coming from a deliberate application of that theory. This brings us to the second similarity between epistemic intuitions and perceptual appearances, a similarity in their relationship to general theories. Just as perceptual observation can conflict with empirical theory, epistemic intuition can conflict with epistemological theory. Whether epistemic intuitions constitute an equally valid independent check on theory is however a point of considerable disagreement. The third way in which the metaphor of appearance figures in discussions of epistemic intuition is in the suggestion that appearance and reality may diverge for intuitions as they do for perceptions: even if we can have no conception of reality without receiving some guidance from sensory appearances, we also know that sometimes things are not as they appear. If sensory appearances may be overridden by background information about the conditions of observation or of one's faculties, or may even be set aside as unexplained anomalies when an otherwise well-established theory is embraced, epistemic intuitions are also in some contexts overridden or judged untrustworthy. Many epistemologists take epistemic intuitions to yield *prima facie* rather than conclusive evidence about the nature of knowledge or justification, subject to correction in the light of a systematic theory, or perhaps subject to outright rejection in cases where there is something problematic about the conditions in which the relevant intuitions have arisen.

There are several ways of understanding the relationship between particular case intuitions and more general theories in epistemology. Advocates of *reflective equilibrium* maintain that philosophical progress is made by adjusting our general theories to better match judgments concerning particular cases while also adjusting our judgments about particular cases to conform to our general theories. In Nelson Goodman's account of the method, we undergo 'dual adjustment between definition and usage, whereby the usage informs the definition, which in turn guides extension of the usage' (66). Applying this method in epistemology, we would start with clear examples of knowledge and produce a definition of knowledge that will help us to classify new or unclear cases; subsequent reflection on these new particular judgments is to help us refine our definition. Reflective equilibrium is sometimes criticized as an unhelpfully conservative exercise of consolidating our existing conceptual practices, but at least in Goodman's version of the procedure it is supposed to include moments of conceptual reform: 'in the interest of convenience or theoretical utility', Goodman suggests, we may 'deliberately permit a

definition to run counter to clear mandates of common usage' (66). In such cases the introduction of the new definition can serve to alter, rather than merely extend, existing usage.

Goodman does not directly characterize reflective equilibrium as an interplay between theory and intuition: what is adjusted in response to the definition is the usage of a term, rather than the particular case intuitions associated with that term. It is consistent with Goodman's description of the process that we may come to use a term to classify particular cases one way while still experiencing the intuitive pull of some other way of judging them. Rudolf Carnap offers the example of the transition from the pre-scientific phenomenal notion of warmth to the modern scientific concept of temperature: even after we appreciate that the latter notion has greater precision and simplicity, the same room will seem to us to be warmer when entered from the snowstorm than from the sauna. Our considered judgments will however come to reflect the scientific meaning of the term: we learn to discount our phenomenal sensations in favor of thermometer readings, and judge that the room has a constant warmth which merely seems to vary.

Carnap is an advocate not of reflective equilibrium but of a method he calls *explication*, in which we refine a messy and vague pre-scientific concept (the 'explicandum') into a simpler and more exact scientific term (the 'explicatum'). A successful explication delivers an exactly defined term that applies to most of the terms once picked out by the explicandum; this new term should be both simple and fruitful, readily connected to an existing network of scientific concepts and helpful in the formulation of new laws. Where reflective equilibrium appears to give equal weight to particular and general judgments, explication assigns them quite different roles. The process of explication starts from examples of current usage, or judgments of particular cases, but once a scientific definition is formed it is not subject to further constraint from reflection on the intuitiveness of its application to particular instances. Instead Carnap contends that following a successful explication our intuitions about particular cases will in many cases eventually come to reflect the new scientific meaning of the term (*Logical Foundations* 13). In other cases we may use the scientific meaning of the term in contexts where precision is required while keeping the original explicandum for everyday use; for example, we may agree to use the term 'or' only in a non-exclusive sense in our logic texts while leaving it ambiguous in ordinary language ('Strawson on Linguistic Naturalism').

One might wonder about the extent to which the adoption of a theory of knowledge would tend to shape one's intuitive judgments concerning particular cases. Intuitive judgments in other areas do seem to be sensitive to the adoption of theories. For example, professional linguists come to have intuitions of grammatical acceptability that are biased in the direction of their theoretical frameworks, and need to test their theories against the

intuitions of neutral informants (Spencer). If a comparable tendency emerges in epistemology, advocates of reflective equilibrium should be expected to urge epistemologists to check their theories against the particular case judgments of nonexperts, on pain of losing a useful source of rational constraint, while advocates of explication should be expected to defend their theories on the basis of theoretical considerations such as simplicity and fruitfulness, and to anticipate that the successful theory of knowledge that emerges will ultimately either bring renegade particular case intuitions into line, or enable us to dismiss them as irrelevant to the philosophically important and exact meaning of 'knowledge' that has been produced by epistemologists.

Reflective equilibrium and explication are two broad ways in which intuitions might relate to theories in epistemology, but they are not the only ways. One might dispute the claim that epistemologists should aim to develop definitions or exact scientific concepts to supercede our pre-theoretical notions of knowledge, justification and the like; one might wonder whether epistemic concepts such as knowledge or justification even pick out theoretically significant phenomena or natural kinds. Discussion of these issues will be taken up after a brief overview of the history of the current debate over intuitions.

## 2. *The Origins of the Present Controversy over Epistemic Intuitions*

As the example from the *Theaetetus* shows, the device of testing theories of knowledge against evaluations of particular cases is not a new one, but new concerns about it have emerged in recent years. The surge of contemporary interest in epistemic intuitions dates to 1963, when Edmund Gettier published a brief but very influential paper which develops two particular cases to generate intuitions cutting against the analysis of knowledge as justified true belief (the JTB theory, for short). In one of the cases, a man who has applied for a job (Smith) has what seems like great evidence that the job will go to his rival (Jones). For some reason Smith decides to count the number of coins in Jones's pocket, and discovers there are ten of them; he then deduces that the man who will get the job has ten coins in his pocket. Meanwhile, notwithstanding Smith's apparently solid evidence to the contrary, the job will in fact go not to Jones but to Smith himself; furthermore, unbeknownst to Smith, Smith himself also has ten coins in his pocket. When Smith believes that the man who will get the job has ten coins in his pocket, Gettier concludes, his belief is true (because Smith himself will get the job) and justified (because of Smith's justified beliefs about Jones); nevertheless, we sense that Smith's belief does not count as an instance of knowledge. The second case in Gettier's article is similar to the first in that the subject relies on a false but justified premise while reasoning responsibly to a conclusion that just happens to be true. However, it was soon established

that reliance on a false premise is not essential to producing a case in which a subject appears to have justified true belief without knowledge (Feldman); indeed, Gettier-type cases can be constructed in a great many other ways.

There is no swift fix to the challenge Gettier raises. Dozens of attempts to construct a better analysis of knowledge emerged in the two decades following the appearance of Gettier's cases. Some added a fourth condition to the original conditions of justification, truth and belief; others aimed to forge connections (such as causal and counterfactual dependence) between two or more of these conditions. After a series of candidate theories of knowledge were rebutted by appeal to increasingly creative and sometimes far-fetched counter-examples (surveyed in Shope), doubts began to emerge about the possibility of generating an analysis of knowledge that would accord with all our intuitive judgments about particular cases. Meanwhile, as counter-examples became more elaborate – involving subjects with strange new perceptual faculties or paranormal powers – it was also found that these cases did not always trigger the sort of robust and widely shared responses that greeted Gettier's original examples, and epistemologists began to wonder which intuitive reactions merit trust, and which should be rejected rather than accommodated.

One way to explain awkward or weak intuitions about remote hypothetical cases was advanced by Alvin Goldman, who proposed that epistemic intuitions are generated when we form a rough impression of the similarities between the cognitive processes at work in the described case and previously stored lists of stereotypical cognitive virtues, like ordinary sight, hearing and logical reasoning, and cognitive vices, like guesswork and wishful thinking. Confusion can arise where some belief-forming process we'd ordinarily regard as a vice is stipulated to be a reliable process within the hypothetical scenario of the case, as in BonJour's thought experiments involving a subject for whom trust in novel psychic powers reliably produces true beliefs (BonJour 1980). If we intuitively judge that such a subject would lack justification, Goldman suggests, this reaction should not be taken as direct evidence against the theory that justification really is a matter of reliability. Because of our background knowledge that trust in the paranormal tends to produce false beliefs in the real world, when we sit down to read the hypothetical case this kind of cognition is already mentally classified as a vice. Goldman thinks that our natural tendency towards 'categorical conservatism' leads us to retain our sense that it is a vice when we first think about its application in the imaginary scenario BonJour has described, notwithstanding the stipulation that this process is reliable in that environment (Goldman 279). However, extensive and careful reflection on an imaginary scenario should be expected to reshape one's sense of what would count as an epistemic vice or virtue in that setting. If our grasp of intellectual virtues is anchored in our past real-world experience through brute psychological

inertia, Goldman reasons, naïve immediate responses to remote hypothetical cases should be interpreted with that fact in mind.

Troubles with intuitions are not restricted to remote hypothetical cases, however. As Gettier cases were becoming increasingly complex, another cluster of problematic intuitions came to light in connection with a debate over whether epistemic vocabulary is context-sensitive, and many of the intuitions elicited in this debate concerned ordinary topics discussed in familiar settings. For example, one would ordinarily be inclined to say that a quick glance at an airline schedule would be enough to enable a person to know that an upcoming flight was going to have a layover in Chicago. If the subject of the case is a person whose life depends on the layover happening as scheduled, however, we might well feel he would need more evidence before we'd be comfortable in judging him to know (Cohen). Differences in the alternatives made salient can also cause epistemic intuitions to waver: for example, just after a discussion of misprints in airline schedules or last-minute changes in flight plans, one might be more reluctant to judge a subject as having knowledge on the basis of a quick glance at a schedule (DeRose; Lewis, 'Elusive Knowledge').

How to account for these shifts remains a topic of considerable controversy. Advocates of subject-sensitive or interest-relative invariantism (IRI) argue that a subject's epistemic position is partly determined by what is at stake for him. According to IRI, the practical interests of the subject matter in determining whether or not he has knowledge, so that the person who has an urgent interest in the layover needs more evidence before he can count as knowing (Hawthorne; Stanley). Sticking to the more traditional view that whether one has knowledge is determined only by truth-conducive factors such as the amount of evidence one possesses, the rival contextualist theory focuses not on the subject but on the *ascriber* of knowledge, and argue that ascribers use 'know' to refer to different (perhaps more or less demanding) relations in different settings. Contextualists maintain that when stakes are high or there are concerns about error, ascribers use the word 'know' to denote a more stringent relation; in more casual contexts 'know' picks out a relation that is more lax (DeRose; Lewis, 'Elusive Knowledge'; Cohen). Meanwhile, strict invariantists tend to argue that we are making a mistake in shifting from ascribing to denying knowledge when factors such as the subject's practical interests are mentioned, for example, because we are confused about the difference between knowing and knowing that one knows (Williamson, 'Contextualism'). The question of whether and how epistemic vocabulary is context sensitive is in any event a difficult and subtle question (Ludlow).

Debates over context-sensitivity uncovered a closely related class of epistemic intuitions that are hard to reconcile with one another. It seems intuitively acceptable for a student with limited financial resources to claim to know that she will not have enough money to buy a mansion this year. It also seems intuitively acceptable for her to claim that she does

not know that her ticket in this week's super-lotto will not win. But if, as seems reasonable, we credit her with the knowledge that a super-lotto win would yield enough money to buy a mansion, then we find ourselves with a paradoxical pattern of epistemic intuitions. If she really does *know* that she will not be able to afford the mansion, she should be able to *deduce* that her lottery ticket will not win from what she knows about what she can afford. If deduction is a safe way of gaining knowledge – and it is hard to think of any ways that are safer – why do we retain the intuition that she does not know that her ticket will lose? Discussed by philosophers like Gilbert Harman in connection with lotteries, this problematic pattern of intuitions appears across a number of other domains where knowledge of a remote possibility seems barred to us, but an assumption about that remote possibility nevertheless seems to be entailed by some mundane claim we would ordinarily take ourselves to know (Vogel). The judgments generating the paradox are each intuitively acceptable, but it is hard to see how any theory of knowledge could render them compatible. Discussing the counterintuitive consequences of his own treatment of the problem, John Hawthorne comments, 'As far as I can see, every candidate story about our puzzle has counterintuitive results' (162).

A final source of worry about epistemic intuitions comes from some studies of the epistemic intuitions of subjects of different cultural backgrounds and socio-economic status (SES). Jonathan Weinberg, Shaun Nichols, and Stephen Stich tested members of different groups on cases such as the following:

Bob has a friend, Jill, who has driven a Buick for many years. Bob therefore thinks that Jill drives an American car. He is not aware, however, that her Buick has recently been stolen, and he is also not aware that Jill has replaced it with a Pontiac, which is a different kind of American car. Does Bob really know that Jill drives an American car, or does he only believe it? (443)

Subjects had to circle either 'Really Knows' or 'Only Believes' after reading the story. In a population of Rutgers undergraduates, only 26% of those who had identified themselves as culturally Western saw Bob as having knowledge here, but 57% of East Asian participants and 61% of Indian Subcontinental participants circled 'Really Knows'. These statistically significant differences in epistemic intuitions could come as a surprise to an epistemologist who takes this story to be an obvious Gettier case, expecting everyone to share his intuition that Bob doesn't know, but only believes.

Weinberg, Nichols, and Stich also found systematic differences between the reactions of different cultural groups to scenarios in which either a single individual or a whole tribe undergoes a freak accident which results in the acquisition of a new quasi-perceptual faculty. In these variations on Keith Lehrer's famous 'Truetemp' case, the subject comes to have a



perfectly accurate sense of the ambient temperature, but has no reason to believe that his new faculty is reliable. Most participants in the experiments judged the subject who relies on this novel faculty to have mere belief rather than knowledge concerning the temperature, but East Asian subjects were more likely than their Western peers to judge the isolated individual as lacking knowledge. Where an entire community was described as having the new faculty, however, the difference was reversed, and East Asian subjects were more likely than their Western peers to ascribe knowledge to a community member using the novel faculty. Weinberg, Nichols, and Stich see these variations as tied to larger cultural differences between the groups tested: specifically, to a pattern of differences studied by social psychologist Richard Nisbett and colleagues, in which an ‘individualistic’ Western emphasis on analyzing objects in isolation from context is contrasted with a ‘holistic’ East Asian emphasis on relationships between objects and their context.

Epistemic intuitions were also found to differ between high- and low-SES groups. After sorting adults into one category or the other by asking whether they had attended college at any point (a standard way of assessing socio-economic status in social psychology experiments), Weinberg, Nichols, and Stich discovered that mentioning a non-actual possibility of error triggers a more critical response among high-SES subjects than their low-SES counterparts. Do high-SES subjects have a higher standard for what counts as knowledge? Are their epistemic intuitions more sensitive to non-actual possibilities for some other reason?

Whatever the answer to such questions, Weinberg, Nichols, and Stich take the apparent cultural and SES variation in epistemic intuitions to constitute a serious problem for the method they see as driving much analytic epistemology, which they dub ‘Intuition-Driven Romanticism’ (IDR). As they characterize it, IDR assumes that universal knowledge of the right epistemic norms is buried within all of us, waiting to be brought to the surface in the form of reactions to particular cases. Given inputs of epistemic intuition, IDR aims to derive universally valid rules about how we should form our own beliefs and assess the beliefs of others. To count as a practitioner of IDR one need not make epistemic intuition the *only* source of input or data for one’s epistemological theories, but one must take intuition to have at least the following weight in one’s theorizing: an input of significantly different epistemic intuitions should result in a significantly different epistemological theory. If quite different epistemic intuitions are found in different cultures, however, then it is not clear how these differing reactions to cases can reveal the nature of a single set of universally valid epistemic norms. According to Weinberg, Nichols, and Stich, the philosopher who probes particular case intuitions is really doing ‘ethno-epistemology’, or studying the contingent epistemic attitudes of the subculture he is probing; as a result, they conclude, intuition-driven

epistemology does not give us normative guidance about how all people ought to think.

Several questions have been raised about the interpretation of this data. Ernest Sosa has argued that it is not clear that variation in the responses to the Bob and Jill Gettier case are properly ascribed to cultural variations in the epistemic realm rather than, say, cultural variations in way readers would flesh out the non-epistemic background details in the stories presented. How likely is it that someone who typically drives an American car is still driving an American car? This is the sort of topic on which subjects of different cultural origin might have different beliefs (Sosa). Linguistic background might also make a difference to how participants construed their task: when asked whether Bob knows that Jill *drives* an American car, it makes a difference whether you understand this as a question about what she typically tends to drive, or as a question about what she is driving right now, since her latest car purchase. For a participant who speaks a language in which the progressive and the recurring forms of the present tense tend to be clearly differentiated, the natural reading of this question might differ from the natural reading for monolingual Western participants. In any event, further experiments with clearer and more detailed sets of cases could shed light on the involvement of such factors. Further study could also make it clearer whether responses to complex hypothetical stories really differ between high and low-SES participants because of variations in their underlying concepts of knowledge and justification rather than, say, variations in verbal reasoning skills or some other factors.

The differences in reactions to the Truetemp cases raise some deeper questions about what should count as an epistemic rule. If *thinking like your peers* is valued more highly or seen to be more reliable in one culture than another, does this show that the two cultures have different epistemic rules? The defender of Intuition-Driven Romanticism could argue that a single set of more fundamental epistemic rules could still be in force across all cultures: if one fundamental rule demands reliability, for example, and reliability is for social or historical reasons associated with conformity in one culture, and with individual initiative in another, then differing responses to the Truetemp cases do not show that epistemic intuitions lack cross-cultural validity at the fundamental level. Fundamental rules like '*seek reliability*' could be the same while local interpretations of these rules differ. Making such a move requires undertaking the doubtless quite difficult task of sorting out the fundamental rules from their contingent local interpretations; it might also require seeing some recent debates in epistemology as taking place at the superficial level of interpretation, rather than the deeper level of norms. Weinberg, Nichols, and Stich are themselves open to the notion that some norms are universally held: having discovered that some questions produce the same reactions across all the cultural and SES groups they have tested, they suggest that 'there

may well be a universal core' to folk epistemology. For example, virtually everyone in all groups agreed that feeling very sure that the next coin toss will come up heads does not mean that you know it will.

Some further empirical data relevant to the 'universal core' hypothesis will be reviewed in section 4; for a more detailed survey of empirical work on variation in epistemic intuitions, and the philosophical significance of this variation, see (Alexander and Weinberg). Meanwhile, any epistemologist who uses epistemic intuitions as a source of evidence can benefit from the cautionary message conveyed by these studies of variation: reactions to sketchy cases can involve a complex array of factors, and one should not be hasty to assume that one's own initial reactions are always definitive.

### 3. Recent Philosophical Views of Epistemic Intuition

On the topic of epistemic intuition, philosophers can be grouped into three broad categories: the defenders, the attackers, and a neutral group we'll call the minimalists. This section surveys a few prominent members of each camp.

Among the defenders, Alvin Goldman and Joel Pust advocate a view of intuitions that they call *mentalism*. According to mentalism, epistemic intuitions are a basic source of evidence concerning our epistemic concepts, where a concept is understood to be a psychologically real state that enables a person to use a natural-language word like 'knows' or 'justified' (187–8). Because my concept of knowledge is what enables me to say, for example, that the subject of a clear Gettier case does not know, this intuitive response automatically reveals something about the shape of that concept. This is not to say that every epistemic intuition delivers direct evidence about the epistemic concepts of the person judging: there are some conditions under which case reactions can be misleading. Calling epistemic intuitions a basic source of evidence amounts to putting them on a par with visual or auditory seemings; these states are *usually* reliable indicators of the truth of their contents. Where poor lighting or heavy background noise might compromise the deliverances of the senses, the distorting factor Goldman and Pust identify as most compromising to epistemic intuition is theory contamination. One's explicit theory of knowledge may be at odds with one's underlying concept of knowledge, as we saw in the initial example involving Theaetetus. If I am a philosopher who is more dogmatically devoted to some philosophical theory than Theaetetus was, then there is a risk that my responses to particular cases may come to be generated by this theory rather than by my underlying concept of knowledge. If the underlying concept is the prime target of philosophical analysis, then the epistemic intuitions of people strongly committed to theories will be less informative than the epistemic intuitions of ordinary folk. According to Goldman and Pust, 'philosophers

rightly prefer informants who can provide *pre-theoretical* intuitions about the targets of philosophical analysis, rather than informants who have a theoretical “stake” or “axe to grind” (183). This is not to say Goldman and Pust think all pre-theoretical intuitions are the same: they are open to the idea that there might be some significant variation in the epistemic concepts of people from different cultural groups.

One might be concerned that the Goldman and Pust program restricts epistemology to the conservative job of describing our existing concepts as opposed to revising and improving them. Goldman and Pust argue that philosophical analysis reveals the present shape of our concepts, but note that philosophy itself can go beyond philosophical analysis: the epistemologist is free to criticize our folk epistemological categories, to expose confusions inherent in them and propose remedies, although in this type of philosophical work ‘intuitions are less likely to play so prominent a role’ (191).

Another defender of intuitions who places a strong emphasis on their relationship with our underlying concepts is George Bealer; for Bealer, however, in excavating our underlying concepts we are not engaged in the naturalistic task of articulating our folk epistemic categories, but in the distinctively philosophical project of uncovering the truth about knowledge, justification, and related notions. Proper philosophical intuition not only shows us how we currently apply a term like ‘knowledge’; it shows us how we ought to do so. Bealer takes pains to distinguish the kind of state he regards as a proper philosophical intuition from other states that are sometimes referred to as intuitive. One of the chief characteristics of philosophical intuition is the appearance of necessity: roughly, ‘if  $x$  intuits that  $P$ , it seems to  $x$  that  $P$  and also that necessarily  $P$ ’ (207). This state is contrasted to, say, the way it appears to us that a house will collapse if its foundations are undermined; ‘plainly, it is *possible* for a house undermined to remain in its original position, or, indeed, to rise up’ (207). Philosophical intuition is also distinguished from belief. Some propositions seem to be necessary but are not believed: Bealer’s standard example is the naïve comprehension axiom (the principle that any predicate defines a set, a principle that naively seems right but leads to paradox). Many propositions are believed but without the sort of immediate seeming characteristic of intuition; for example, remembered historical or mathematical facts, or the products of complex calculations. It doesn’t just immediately *seem* to you that  $869 + 544 = 1413$ , even if you believe this is so. In order to qualify as an epistemic intuition for Bealer, a cognitive seeming must be immediate and fairly robust, both in the sense of resisting change in the face of countervailing beliefs (so the naïve comprehension axiom can still *seem* right even when we no longer believe that it is), and in the sense of resisting change in different contexts (if a Gettier case response is a proper intuition, it should strike us the same way whenever we read the case). Bealer suggests that results like those of

Weinberg, Stich, and Nichols are best explained by a failure to isolate the right kind of robust intuition for study (213).

Bealer grants that intuitions are locally fallible, as the set theory paradoxes show, but he insists that the large-scale project of working with intuitions to systematize our philosophical knowledge is not wrecked by this local fallibility. He argues that for philosophical concepts of the type he is probing it is constitutive of concept possession that efforts to systematize our intuitions should over time and under appropriate cognitive conditions lead us to increasingly stable theories that converge on the truth. We possess a concept *determinately*, according to Bealer, if we have natural propositional attitudes (such as belief) towards propositions in which this concept figures, and have these attitudes not just on the basis of some partial or inaccurate understanding of the content of this concept (222). Bealer contends that it is possible for us to gain determinate possession of the core notions of epistemology (knowledge, justification, belief), and that our natural epistemic intuitions form an indispensable basis for attaining this aim. Without reliance on intuition we would have no basic evidence for our philosophical claims, and would find ourselves sliding into a regressive search for evidence for them. According to Bealer, even philosophical programs that claim to be skeptical about intuitions are ultimately dependent on intuitions for their own support, a point also pressed by Joel Pust in his 'Against Explanationist Skepticism'.

Attacks on intuitions come from a variety of angles. Hilary Kornblith allows that Bealer, Goldman, and Pust may be right to claim that a person's epistemic intuitions are a valuable source of evidence about her concept of knowledge; however, Kornblith draws a sharp distinction between the project of finding out about our concepts of knowledge and the project of finding out about knowledge itself ('Naturalism and Intuitions'). Just as metallurgy makes progress by examining the properties of aluminium itself, rather than the properties of pre-theoretical concepts of aluminium, so also epistemology should examine knowledge rather than pre-theoretical attitudes to it. Pre-theoretical concepts of knowledge are a legitimate object of study for the sociologist or historian; investigating these folk concepts in an accurate and responsible manner requires the collection of extensive empirical data and awareness of some basic psychological facts about concepts. Kornblith highlights recent empirical work on folk concepts that suggests that they are not usually represented in us as sets of necessary and sufficient conditions, but rather by means of prototypical examples together with a multi-dimensional range of ways and degrees to which members of the category can deviate from the prototype. Folk concepts like 'game', 'vehicle', and arguably 'knowledge' do not exhibit the sort of necessary-and-sufficient-conditions unity that analytic philosophers are looking for, Kornblith contends. Knowledge itself might exhibit such unity, however, especially if it is, as he argues elsewhere, a natural kind (*Knowledge and its Place in Nature*).

Kornblith contends that routine philosophical practice shows some awareness of the problems with pre-theoretical folk concepts. Although professional philosophers solicit epistemic intuitions from beginners in their courses, they are swift to dismiss those intuitions that do not fit the central accepted principles of current epistemology. When beginning students claim to find it intuitive that knowledge can be false (because 'what the experts know now might turn out to be wrong'), or that what is known need not be believed (because 'after studying boring material for a test, I might end up knowing lots of stuff without believing any of it'), their remarks are usually set aside as unconnected to the concept of knowledge that counts. Such intuitions are indeed difficult to unify into any coherent theory, but if philosophers really were interested in the folk concept of knowledge they would have to take these intuitions at face value, Kornblith contends, recognizing that folk concepts can be scattered and messy things. While they throw out the ideas of beginners, philosophers are keen to look at their own intuitions and those of their colleagues, Kornblith notes. If Goldman and Pust are right that theoretical contamination is the largest risk factor for distorted intuition, then in soliciting intuitions from those who dedicate themselves to developing theories of the target concepts, philosophers have chosen the worst possible population to poll. Kornblith thinks that epistemologists are in fact right to look to the judgments of informed investigators in the course of improving their theories, exactly because the goal of epistemology is not to figure out folk concepts but to figure out the character of knowledge itself. To the extent that professional epistemologists have unusually systematic and coherent sets of responses to epistemic scenarios, these responses should indeed tell us less about the pre-theoretical concept of knowledge, but this does not mean that the responses of professionals must tell us less about knowledge itself, as long as knowledge itself is a natural kind amenable to systematic theorizing. Leaving the study of folk concepts to sociologists, Kornblith himself thinks that the best route to figuring out the character of knowledge itself is to study its evolutionary role in guiding successful action.

Where Kornblith counsels epistemologists to turn their backs on intuitions altogether, Brian Weatherson advises a more selective attack on certain problematic intuitions, balanced with respect for other, preferably stronger intuitions, in a manner motivated by a number of theoretical considerations. In Weatherson's view, 'the true theory of knowledge is the one that does best at (a) accounting for as many as possible of our intuitions about knowledge while (b) remaining systematic' (7). Weatherson notes that the classical view of knowledge as justified true belief is appealingly systematic: it is a simple analysis of knowledge into theoretically significant terms, a view whose adoption would make it an appropriately non-arbitrary matter whether a given state should count as one of knowing. While the JTB theory violates intuition in Gettier cases, there is enough

internal conflict among commonly held intuitions concerning knowledge that some violation of intuition appears to be inevitable. Given the theoretical virtues of the JTB theory, Weatherson argues, we might do well to reject our Gettier-case intuitions.

Weatherson is happy to grant that we do have a disposition to use the word 'knowledge' in ways that go against the JTB theory: it certainly seems to us that Smith does not know that the man who will get the job has ten coins in his pocket. However, the meanings of our terms are not just set by our existing dispositions to use those terms: when the medievals classified Venus as a star or whales as fish, they were making mistakes about planets and mammals rather than just meaning something different by 'star' and 'fish'. Both for the medievals and for us, meaningful predicates like 'star' and 'fish' denote natural properties, properties that can figure in simple and systematic theories. To shift our usage so that whales are not called fish might have appeared counter-intuitive to our ancestors but was something we needed to do, given the real or natural meaning of the term. Embracing David Lewis's view that the meaning of a predicate is the natural property that falls closest to our disposition to use that predicate (Lewis, 'New Work'), Weatherson contends that the meaning of 'knowledge' must be somewhat out of line with our existing dispositions to use it. Given our disposition to see the subjects of Gettier cases as lacking knowledge, and the tremendous difficulty we have had in coming up with a simple and systematic theory that accommodates those tricky intuitions, Weatherson concludes that 'there just are no reasonably natural properties in the neighborhood of our disposition to use 'knows' (13). The best we can do is to venture outside of that immediate neighborhood and search for the closest available contender, some simple theory that preserves as many of our intuitions about knowledge as possible. Weatherson suggests the best candidate for this job may be the JTB theory, and that epistemologists should keep in mind the possibility that Gettier intuitions provide only apparent rather than decisive evidence against this classical view.

Anyone who sets out to reject our Gettier intuitions needs to explain how intuitions can be untrustworthy. Weatherson cites Robert Cummins's argument against trust in philosophical intuition: according to Cummins, epistemic intuition is a poor guide to the truth about knowledge because intuition cannot be calibrated. Intuition differs from perceptual observation in this respect: we can calibrate or double-check the deliverances of vision against touch, or the deliverances of telescopically aided vision against what the naked eye can reveal, say by examining distant mountains through the telescope and then looking at them from close up. In calibrating an information source, we gain independent evidence that it is reliable and also learn something about the range of circumstances under which it is reliable. If pre-theoretical epistemic intuitions are just raw responses to cases, then it is not obvious how they could be calibrated,

Cummins argues: we could discover that many people shared similar intuitions to certain cases, but on its own this finding would not validate these intuitions. Almost everyone presented with the Müller-Lyer illusion (in which parallel lines of equal length are capped by inward and outward-facing arrows) will see the lines as differing in length when they are not; there is not always safety in numbers. We'll return in section 4 to the issue of whether epistemic intuitions can be calibrated; if Cummins and Weatherston are right that they cannot, then they arguably have a good case against naïve trust in intuitions that conflict with the best available theory of knowledge.

The last camp to discuss are the minimalists, who set out neither to attack nor to defend a category of judgments labeled as 'intuitions', but to challenge the notion that there is a special category of judgments worthy of this name, a group of reactions or responses with special evidential force, different in fundamental character from other judgments we might make, for example, on the basis of inference.

The word 'intuition' does not appear in Gettier's original article. Jaakko Hintikka notes that the explicit appeal to states labeled as intuitions becomes popular in philosophy only after the mid-1960s (127). He contends that this device began as an effort to make the methods of philosophy resemble the newly successful methods of Noam Chomsky's linguistic theory. Chomsky developed a powerful theory of an innate universal grammar by appeal to the intuitions of speakers of a natural language (*Syntactic Structures*; *Cartesian Linguistics*); presenting itself as intuition-driven was one way in which philosophy might continue to appear relevant and up-to-date as some of its traditional topics fell under the scope of empirical investigation. Hintikka then argues that this effort to appropriate the methods of linguistics for philosophy is a bad idea. When linguists probe the human language faculty by eliciting intuitions of acceptability from competent speakers, there is no standard outside the language faculty itself against which these intuitions could be found wanting. Philosophical intuitions, on the other hand, typically purport to tell us about objective phenomena like the nature of knowledge or justification, and the distance between our intuitive responses and the phenomena themselves may be more problematic. In addition, Chomsky's appeal to intuitions made sense in the context of a theory which insisted that our linguistic capacities had an innate underlying basis – Hintikka notes that Chomsky's own justification for his early stance on intuitions was quite explicitly Cartesian (132–3) – but philosophers who doubt the reality of innate philosophical ideas (that is, the majority of philosophers) have no parallel basis for their appeals to intuition. Arguing that philosophers have been guilty of 'reifications and mystifications' of intuitions, Hintikka contends that 'what are explicitly called intuitions often turn out to be products of perfectly ordinary discursive thinking, combined with suitable observations' (143). Drawing attention to the fallibility of our



spontaneous judgments, Hintikka urges that it is no defense of a claim to label it 'an intuition', and suggests (half-jokingly, he says) that no philosophy papers appealing to intuitions should be published anywhere, 'unless the basis of those appeals is made explicit' (147).

More recently Timothy Williamson has also criticized the view that 'intuitions' form a distinct category of mental states with special justificatory force; rather, Williamson argues, 'what are ordinarily called 'intuitions' in philosophy are just applications of our ordinary capacities for judgment' (109). His diagnosis of the intuition trend is different from Hintikka's, however. Williamson maintains that the drive to label these judgments 'intuitions' comes from a desire to find a firmly uncontroversial and completely identifiable evidential basis for philosophy. Having read Weatherson, we realize that it is philosophically controversial whether the subject of a Gettier case does or does not have knowledge; we might try to take comfort in the thought that it is less controversial whether a given person does or does not have the epistemic intuition that the subject of a Gettier case has knowledge. If we want to know about what intuitions people actually have, we can use familiar empirical methods to find out: we can conduct polls and determine that, say, 74% of self-identified Western Rutgers undergraduates have the intuition that Bob does not know that Jill drives an American car. However, Williamson argues, even empirical facts of this sort are not given to us with perfect transparency; especially if we consider questions about, say, the relative strengths of two subtly different epistemic intuitions we might have, we are also fallible concerning our epistemic intuitions ('Philosophical Intuitions and Scepticism' 120–1). The move from talking about knowledge to talking about intuitions of knowledge is not a move that insulates us from the possibility of error.

Furthermore, Williamson maintains, there are serious costs to construing our evidence as consisting in awareness of our own psychological states. If, instead of being able to say directly that *the victim of Gettier case lacks knowledge* we must instead see our basic evidence as being nothing more than *I have an epistemic intuition to the effect that the victim of the Gettier case lacks knowledge*, then according to Williamson we have opened up a skeptical gap between our judgments and the phenomena, a gap that will be extremely hard to close. If we would like to model our methods on those of the sciences, Williamson contends, we should recognize that empirical investigators do not hedge their evidence claims in the way proposed by advocates of a distinct evidential category of intuitions. A scientist investigating the relationship between atmospheric pressure and boiling point says not just that *he feels inclined to judge that the barometer reading is 101 kPa* but that *the pressure is 101 kPa*. According to Williamson, there is no reason why we should not consider ourselves to be directly in touch with the phenomena in epistemology as well: a judgment concerning a particular case is not simply a source of evidence about our own

psychological states or concepts, but a source of evidence about knowledge itself. Finding that there is evidence of variation in epistemic intuitions (as in the Weinberg, Stich, and Nichols experiments) should also be handled as we handle similar findings concerning variation in ordinary observational evidence, for example, in eyewitness testimony. Discovering that eyewitness reports are often inaccurate should trigger not general skepticism about sense perception but specific forms of caution about the information conveyed by these reports under various conditions; discovering that intuitions can vary among subjects should likewise prompt caution rather than wholesale rejection, Williamson contends ('Philosophical Intuitions and Scepticism' 150).

Even if Williamson is right that epistemology should not be taken to rest on a class of uncontroversial judgments with special evidential value, it could still be true that the states known to us as epistemic intuitions have some interesting common features, and that knowledge of these features can improve our capacity to make accurate inferences from them. With this in mind, the next section examines some empirical research on epistemic intuitions.

#### 4. *Empirical Evidence from Linguistics and Psychology*

In recent years psychologists have devoted considerable attention to the processes that enable us to assess our own epistemic states and those of other people. It is not assumed that exactly the same processes support self-monitoring and the ascription of mental states to others; indeed many psychologists study one of these functions to the exclusion of the other. Because most epistemological case intuitions involve assessments of others (like Plato's jury and Gettier's Mr Smith), this survey will concentrate on empirical work on third-person assessments, but a brief treatment of self-assessment is also in order. It is worth noting that intuitive self-assessments do figure in epistemological work, sometimes quite prominently. For example, David Lewis's 'Elusive Knowledge' takes a number of first-person knowledge claims as its point of departure, and it is not uncommon for epistemologists to encourage us to reflect on various beliefs of our own to sense the intuitive pull of classifying beliefs of a given type as clear cases of knowledge.

Much work on intuitive epistemic self-assessment has focused on a phenomenal state that has come to be known as the *feeling of knowing* (or FOK) state. The FOK state is what one feels when one senses that one can retrieve an answer to a question even prior to actual retrieval from memory, a state which is prolonged and felt most acutely when retrieval is blocked but the answer is 'at the tip of the tongue'. This vivid and protracted FOK state is experienced when a person feels the urge to say '*I know that I know this one*' to a trivia question he cannot immediately answer, but we are thought to experience a more rapid FOK even for

questions we can answer directly. This state has predictive value: a strong feeling of knowing on a free recall (fill-in-the-blank) question is correlated with superior ability to choose the objectively correct answer in a subsequent multiple-choice test (Hart). The FOK state is stable, in the sense that it tends to be felt to the same degree for a given question when this question is presented on multiple occasions (Costermans et al.). There is some debate about the function that is served by one's having a conscious feeling that one has knowledge on a given point, but one common suggestion is that the FOK informs our decisions about how to allocate our cognitive resources (e.g., Koriat, 'Feeling of Knowing'). For example, when one has a strong FOK for a given problem on which retrieval is blocked, one will spend much longer searching memory for the answer, rather than attempting some other strategy to try to solve the problem; on the other hand, when a strong FOK accompanies successful retrieval, one is less likely to devote time and energy to double-checking one's answer. Because it is at least roughly accurate, the FOK signal generally improves our decisions about where to invest our cognitive energy, and because this signal is available to consciousness, these decisions can be integrated into conscious deliberations about the pursuit of our goals.

While it was initially thought that the FOK arose from some indirect access to underlying memory traces just out of reach of retrieval, later work revealed that the FOK involves more complex and oblique processes. The feeling of knowing can be generated by familiarity with terms in the question, independently of any present or prior access to an answer (Reder); indeed, the speed of the feeling of knowing is what enables game show contestants to hit the buzzer for questions they are likely to get right even before there has been time for retrieval. The FOK state can also be triggered by the volume of associated information brought to mind by a cue even when this information is not relevant to answering the question (Koriat, 'How Do We Know that We Know?'); in fact, we can be manipulated into having strong FOK states for questions whose answers are wholly unknown to us (Schwartz). Cognitive experiences, like perceptual experiences, are in some cases illusory.

Although it may feel like a moment of 'direct perception' of the quality of our cognition, the conscious FOK state appears to be generated by a unconscious heuristic process that monitors several aspects of memory performance, including speed and amount of recall. Because these factors are only indirectly connected with knowing, the FOK has limited validity as an indication of whether we actually know something, however it may feel to us (Koriat, 'Dissociating Knowing'). It is a live question whether FOK states influence epistemologists' introspective intuitions about whether certain kinds of beliefs constitute knowledge. If the FOK influences (or worse, if it constitutes) epistemologists' intuitive introspective assessments of their own knowledge, then epistemologists should be aware of the conditions under which FOK states can be misleading, and should

anticipate that a good theory of knowledge will yield final judgments about certain cases that are at odds with our intuitive introspective assessments.

There have been some efforts to demonstrate that we have feelings of knowing not only for ourselves, but also for other people. Susan Brennan and Maurice Williams coined the expression ‘feeling of another’s knowing’, or FOAK state, for the feeling that accompanies the rapid assessment we make of the knowledge of others. Rising intonation and certain sorts of pauses decrease the FOAK others feel when you speak; grimaces and avoidance of eye contact are also damaging to FOAK. Like philosophers who have argued that we need some ‘fast and frugal heuristics’ to evaluate what others believe (Spicer), Brennan and Williams believe that our fluctuating FOAK states monitor the utterances of others in a way that enables us to respond selectively to their testimony. Again, these feelings are heuristically based and only indirectly correlated with knowledge itself. It appears that we associate the traits that produce higher FOAK states with higher speaker confidence in what is said; higher speaker confidence often accompanies knowledge. Some limits of the validity of the FOAK state are easy to spot: we recognize on reflection that a person can have high confidence in a proposition without having knowledge of it, or show overt signs of low confidence despite having knowledge. However, recognizing these limits on reflection does not imply that we will cease to feel variation in FOAK states when we encounter speakers who show outward signs of higher or lower confidence. Where stories intended to generate particular case intuitions accidentally include mention of such outward signs of high or low confidence, it is possible that FOAK states could bias our ensuing epistemic intuitions. In composing test stories to elicit epistemic intuitions, epistemologists should be aware of the possibility of FOAK-based distortion, and should be alert to the consequences of describing their subjects in ways that the intended audience would associate with high or low confidence.

The main body of psychological work on epistemic assessments of others is concerned not with the phenomenal FOAK state but with a broader set of cognitive capacities that enable us to represent and evaluate the beliefs of others. These knowledge- and belief-ascribing capacities function quite differently in young children, older children, and adults; indeed, developmental work on the way these natural epistemic capacities mature seems relevant to the question of whether epistemic intuitions are subject to calibration. Cross-cultural commonalities in the development of epistemic evaluations also have some bearing on the question of the extent to which we should expect epistemic intuitions to be universal.

For example, one feature of epistemic evaluation that has appeared in every culture studied is a sharp progression in young children’s capacity to see others as having false beliefs. Emergence of this capacity has been

studied in a great range of settings, from urban daycare centers in New York and Tokyo to indigenous non-literate communities in Peru and Cameroon. Across all cultural groups, the ability to grasp the notion of false belief has been found to emerge between the ages of three and five (Wellman, Cross, and Watson). In the classic false belief test, a treat is hidden in a box in the presence of the child and a third party ('Max'). Max then leaves the room, and the treat is moved to another location in the sight of the child, say, a drawer. The child is then asked where Max will look for the treat when he re-enters the room. Overwhelmingly, three-year-olds fail this test: they say that Max will look in the drawer, where the child herself knows the treat is now located. Five-year-olds pass: they keep track not only of the present location of the treat, but also of Max's (now false) belief (Wimmer and Perner). The switch from fail to pass is not the product of a general early incapacity to think about the minds of others – even at the age of eighteen months, a child can readily grasp the notion that others have desires different from those of the child (Repacholi and Gopnik) – but a more specific development in the understanding of belief. Very young children seem to assume that the world is transparent to others, and have a strong tendency to over-ascribe knowledge (Mossler et al.); older children can keep track of the ways in which others are mistaken, and in predicting the actions of others are able to make use of the concept of a belief as a state which may or may not be an accurate reflection of reality.

The switch to seeing others as capable of misrepresentation is not the only sharp and cross-culturally uniform progression in epistemic evaluation that psychologists have found. Other developmental stages involve increasing refinements about the sources of knowledge: for example, where small children assume that seeing an object gives a person knowledge of all its properties known to the child – including, say, which of two identical-looking cups contains warm water and which contains cold water – older children and adults know that sight alone does not do this for us, and restrict their ascriptions of knowledge accordingly (O'Neill et al.). The ability to keep track of the evidential sources of one's own beliefs – for example, being able to tell whether one saw an event or was only told about it – also rises sharply between the ages of three and five (Woolley and Bruell).

Further developmental changes in our patterns of epistemic intuitions arise with the emergence of the ability to recognize the difference between ignorance and false belief. Children who can pass the false belief test do not yet possess the mature concepts of knowledge and ignorance: for example, they appear to equate *failure to know* with *having a false belief*. In an experiment of Ted Ruffman's, a child and a puppet observer see a dish of red and green candies, and then the puppet leaves. The child sees a red candy taken from the mixed dish and placed in an opaque box, which is then closed. The child is then asked what color of candy the

puppet will think is in the box when he returns and is told that a candy has been moved from the dish to the box. The youngest children answered 'red', unable to form the notion of a false belief. By the age of seven, children answer as adults would ('he won't know whether it's red or green'). Interestingly, children in the middle range had a strong tendency to report that the puppet would think the candy was *green*, suggesting that we pass through a developmental period in which we equate ignorance and false belief and are unable to represent another's suspension of belief on a question whose answer is known to us.

There are cross-cultural differences in the order of acquisition of some of the components of the mature concept of knowledge: for example, American children tend to master 'diverse beliefs' tasks before they master 'knowledge-ignorance' tasks, while Chinese children master these stages in the opposite order (Wellman et al.). (Diverse belief tasks require the child to ascribe different beliefs to different subjects concerning a question whose answer is unknown to the child; knowledge-ignorance tasks require the child to ascribe ignorance to a subject concerning a question whose answer is known to the child.) However, the major points on the scale appear to be cross-culturally uniform: ability to ascribe diverse desires always precedes ability to ascribe diverse beliefs, and ability to ascribe ignorance always precedes ability to ascribe explicit false belief (Wellman, Cross, and Watson; Wellman et al.).

Studies of both normal and impaired social reasoning tend to support the view that the capacity to attribute mental states to others is, like natural language, modular in character: our minds have a set of mechanisms dedicated to mental state ascription and evaluation. Using brain imaging technology, Rebecca Saxe and Nancy Kanwisher have found that certain regions in brain, most notably in the right temporal-parietal junction, are activated when adults read short stories about the beliefs of others, like Plato's story about the jury. These patterns of activation are distinctive of epistemic evaluation, and do not occur when subjects read stories about the desires of others, or about their physical appearance, or when subjects read stories about other sorts of representations. For example, these regions of the brain are engaged when subjects read a story about a person with an outdated belief, but not when they read a logically parallel story about a photograph of a scene that has changed since the photograph was taken (Saxe and Kanwisher).

Saxe and Kanwisher found that typical adults respond faster to questions about the outdated belief stories than to questions about the outdated photograph, but there are populations for whom social reasoning does not come so easily. Autistic children have difficulty ascribing mental states to others, while performing relatively well on comparable tasks concerning non-mental representations like photographs (Baron-Cohen). The selective impairments shown in autism are taken as further evidence for the modularity of mental state ascription. Evidence that some special modules

are involved in epistemic calculation also emerges from the pattern of errors and biases typical adults have in evaluating the beliefs and knowledge of others, for example, in our systematic overestimations of the extent to which others know what we know (Nickerson; Saxe). A more detailed understanding of the biases that afflict spontaneous epistemic judgments could assist philosophers wondering which epistemic intuitions to trust.

There are several actively defended theories about the character of the module involved in social cognition, and its relationship to more general mental capacities like attention and working memory. Current research in social cognitive neuroscience examines such topics as the relationship between domain-specific capacities in the epistemic module and general regulatory capacities (Saxe et al.), and the conditions under which epistemic assessments are automatic rather than deliberate (Satpute and Lieberman). Research in this area may improve our understanding of variations in epistemic intuitions such as the variations motivating contextualist theories of knowledge. For example, if we switch between automatic and deliberate epistemic assessment when we move from everyday contexts to high-stakes contexts or self-conscious reasoning, and if the heuristics of automatic assessment differ in their output from the rules of deliberate assessment, then we should not be surprised to see variations in epistemic intuitions in these different conditions.

Linguists also study the acquisition and deployment of mental state terms. Although ‘know’ and ‘think’ are said to be linguistic universals, occurring in every natural language (Wierzbicka), natural languages differ in the manner and extent to which they express distinctions that have epistemological significance. For example, roughly a quarter of the world’s languages have grammaticalized evidentials (Aikhenvald 1). An evidential is a marker of the origin of the speaker’s information, for example, indicating whether the speaker is reporting something she saw herself, figured out from other evidence, or heard about from others. In English we have the option of marking such distinctions lexically, modifying our claims with words like ‘apparently’ to indicate inference or ‘allegedly’ to indicate hearsay, if and when we choose to do so. In languages with grammaticalized evidentials such markings appear as verbal prefixes, suffixes, or particles, and are often mandatory, or mandatory across a large class of assertions. For example, every past-tense sentence in Turkish must contain a verb suffix indicating whether the source of the proposition asserted is direct (e.g., one’s own perception) or indirect (inference, testimony). Some version of this two-way direct/indirect contrast appears in all languages with grammaticalized evidentials; other languages subdivide sources of evidence further, but never into more than five categories in total, where these categories (visual, non-visual sensory, inferred, conjectured, reported) measure the speaker’s experience of his evidence (Aikhenvald ch. 2).

Linguists have noticed that although there are many other differences in source that might have social significance (e.g., revelation, instinct, teachings of an elder), these never appear as grammatical evidential categories; rather, languages draw from the same restricted system of options (Speas). It has also been observed that there is a cross-linguistic hierarchy in evidentials, with direct evidence always at the top and hearsay at the bottom (Willett). Just as it would be misleading in English to say 'John was allegedly at the party last night' if one had not only heard that he was coming but in fact seen him there, so also in languages with grammaticalized evidentials one is expected to use the verb form that reflects the highest level of information one possesses.

Because evidentials express epistemologically significant distinctions, one might expect to see systematic differences in the epistemic assessments of speakers for whom these distinctions are mandatory (like speakers of Turkish or Korean) and speakers for whom these distinctions are optional (like speakers of English or French). More specifically, one might expect Korean speakers to be more adept at keeping track of the type of evidence supporting their utterances than their English-speaking counterparts, or to develop the relevant abilities at an earlier age. Some recent work by Anna Papafragou and colleagues indicates that this is not the case: comparing the performance of Korean- and English-speaking preschoolers on tasks involving recognition of sources of evidence, they found that language did not appear to make a difference to their young subjects' abilities to monitor the sources of their evidence. The two linguistic groups improved uniformly on source monitoring tasks between three and four years of age, even though one might have predicted an advantage for the Korean children, given their exposure to a language in which evidential markings occur in a great number of the adult sentences they hear. Younger Korean children also showed very limited understanding of the meaning of the evidential markings in their language. Papafragou and colleagues take these results to point to 'the independence of language from source monitoring abilities' (292); they also suggest that the psychological capacities involved in source monitoring could appropriately be seen as universal. In any event, the psychological capacities that enable us to track and evaluate the sources of our beliefs are not as sensitive to linguistic differences as a cultural relativist might expect.

Although much empirical work on epistemic assessments focuses on the early development of children's theory of mind, psychologists and linguists have also studied later developments. In a study comparing understanding of mental state verbs like 'know' in third-graders, sixth-graders, and adults, Rachel Falmagne found evidence of systematic ongoing refinement in our grasp of the meaning of these words (Falmagne et al.). Mental state verbs can be organized according to how much confidence they express: on such a ranking, 'was sure' rates higher than 'assumed'. Mental state verbs can also be organized according to their logical properties, like



factivity. A verb is factive when it presupposes the truth of its complement clause: a speaker who says that Thomas *knows* that his wife is cheating on him must be committed to the claim that Thomas's wife is cheating on him. For a non-factive verb like 'thinks' this sort of inference does not hold: it might be true that Thomas *thinks* his wife is cheating when she is not. What Falmagne found was that in assessing reports of the beliefs of others, confidence mattered more to younger subjects and logical properties mattered more to the older ones. From a sentence like 'Carl was sure that Bill had bought a new skateboard' third graders will readily infer that Bill had in fact bought a new skateboard; adults are much less inclined to do this. As we mature, our understanding of the meaning of 'knows' comes to reflect a better grasp of the differences between feeling sure about something and knowing it, and the logical property of factivity comes to play an increasingly important role in our understanding of knowledge.

Some philosophers have argued that epistemic intuitions are untrustworthy because they are uncalibrated; one question raised by the developmental data is whether the development of our understanding of 'knows' as we mature is evidence of calibration. Epistemic intuitions become more accurate as one graduates from the toddler's vision of a world transparent to every mind to the older child's grasp of the ways in which the source of a person's information matters to the quality of his belief, and more accurate still as one gains the adult insight that feeling sure is not enough for knowledge. Such developments can be seen as the product of a changing theory of knowledge that becomes increasingly successful in helping us understand the words and actions of others. Because the theory generates predictions about what people will say and do, and assessments of what people say and do are not necessarily epistemic assessments themselves, we gain some independent check on epistemic assessment. Some early patterns of epistemic intuition, such as the failure to distinguish ignorance and false belief, are unsuccessful and must be abandoned; other patterns, such as the recognition that the channel of information can make a difference to what is known, are successful and come to be retained in our basic cognitive repertoire. The extent and quality of this natural calibration remains open to question, however. If typical adults naturally end up in agreement on some core set of epistemic intuitions – for example, agreeing that someone who has had a correct hunch about a coin toss did not know how the coin would land – this does not on its own imply that adults naturally end up in agreement on enough particular cases to support the construction of an acceptable theory of knowledge, or even that typical adults are getting it right on the points on which they do naturally agree. Further research on the scope and origin of intuitive epistemic agreement should prove useful to anyone who wonders whether epistemology should aim primarily to capture or to correct our natural epistemic intuitions.

### Acknowledgements

Thanks to Jane Friedman, Tamar Szabó Gendler, Ian Hacking, Hilary Kornblith, Gurpreet Rattan, Sherri Roush, and Sergio Tenenbaum for helpful comments.

### Short Biography

Jennifer Nagel is interested in the nature of epistemic justification, and in relationships between first- and third-person epistemic assessments. Her recent work examines what is available to the subject in the course of ordinary belief formation. She has published papers on empiricism, experience, and the Cartesian response to skepticism. Nagel holds a Ph.D. from the University of Pittsburgh, and teaches philosophy at the University of Toronto.

### Note

\* Correspondence address: Department of Philosophy, University of Toronto, 170 South George Street, Toronto, Ontario M5R 2M8, Canada. Email: jennifer.nagel@utoronto.ca.

### Works Cited

- Aikhenvald, Alexandra. *Evidentiality*. New York, NY: Oxford UP, 2004.
- Alexander, Joshua and Jonathan Weinberg. 'Analytic Epistemology and Experimental Philosophy'. *Philosophy Compass* 2.1 (2007): 56–80.
- Baron-Cohen, Simon. *Mindblindness*. Cambridge, MA: MIT Press/Bradford Books, 1995.
- Bealer, George. 'Intuition and the Autonomy of Philosophy'. *Rethinking Intuition: The Psychology Of Intuition And Its Role In Philosophical Inquiry*. Eds. Michael DePaul and William Ramsey. Oxford: Rowman & Littlefield, 1998. 201–39.
- BonJour, Laurence. 'Externalist Theories of Empirical Knowledge'. *Midwest Studies in Philosophy* 5 (1980): 53–73.
- Brennan, Susan and Maurice Williams. 'The Feeling of Another's Knowing—Prosody and Filled Pauses as Cues to Listeners about the Metacognitive States of Speakers'. *Journal of Memory and Language* 34 (1995): 383–98.
- Carnap, Rudolf. *Logical Foundations of Probability*. 2nd ed. Chicago: U of Chicago P, 1962.
- . 'Strawson on Linguistic Naturalism'. *The Philosophy of Rudolf Carnap*. Ed. Paul Arthur Schilpp. LaSalle: Open Court, 1963. 933–40.
- Chomsky, Noam. *Cartesian Linguistics*. New York, NY: Harper, 1966.
- . *Syntactic Structures*. The Hague: Mouton, 1957.
- Cohen, Stewart. 'Contextualism, Skepticism, and the Structure of Reasons'. *Philosophical Perspectives* 13 (1999): 57–89.
- Costermans, J., G. Lories, and C. Ansay. 'Confidence Level and Feeling of Knowing in Question Answering – the Weight of Inferential Processes'. *Journal of Experimental Psychology: Learning, Memory and Cognition* 18 (1992): 142–50.
- Cummins, Robert. 'Reflection on Reflective Equilibrium'. *Rethinking Intuition: The Psychology of Intuition and its Role in Philosophical Inquiry*. EdS. Michael DePaul and William Ramsey. Oxford: Rowman & Littlefield, 1998. 113–127.
- DeRose, Keith. 'Contextualism and Knowledge Attributions'. *Philosophy and Phenomenological Research* 52 (1992): 913–29.

- Falmagne, Rachel, Joanna Gonsalves, and Sarah Bennett-Lau. 'Children's Linguistic Intuitions about Factive Presuppositions'. *Cognitive Development* 9 (1994): 1–22.
- Feldman, Richard. 'An Alleged Defect in Gettier Counter-examples'. *The Australasian Journal of Philosophy* 52 (1974): 68–9.
- Gettier, Edmund. 'Is Justified True Belief Knowledge?'. *Analysis* 23 (1963): 121–3.
- Goldman, Alvin. 'Epistemic Folkways and Scientific Epistemology'. *Philosophical Issues*. Vol. 3, Science and Knowledge (1993): 271–85.
- and Joel Pust. 'Philosophical Theory and Intuitional Evidence'. *Rethinking Intuition: The Psychology Of Intuition And Its Role In Philosophical Inquiry*. Eds. Michael DePaul and William Ramsey. Oxford: Rowman & Littlefield, 1998. 179–97.
- Goodman, Nelson. *Fact, Fiction, and Forecast*. 2nd ed. New York, NY: Bobbs-Merrill, 1965.
- Harman, Gilbert. 'Knowledge, Inference and Explanation'. *American Philosophical Quarterly* 5 (1968): 164–73.
- Hart, J. T. 'Memory and the Feeling-of-Knowing Experience'. *Journal of Educational Psychology* 56 (1965): 208–16.
- Hawthorne, John. *Knowledge and Lotteries*. New York, NY: Oxford UP, 2004.
- Hintikka, Jaakko. 'The Emperor's New Intuitions'. *The Journal of Philosophy* 96.3 (1999): 127–47.
- Kornblith, Hilary. *Knowledge and its Place in Nature*. New York, NY: Oxford UP, 2002.
- . 'Naturalism and Intuitions'. *Grazer Philosophische Studien: Philosophical Knowledge*. Eds. Christian Beyer and Alex Burri. Amsterdam, The Netherlands: Rodopi, 2007. 27–49.
- Koriat, Asher. 'Dissociating Knowing and the Feeling of Knowing: Further Evidence for the Accessibility Model'. *Journal of Experimental Psychology: General* 124 (1995): 311–33.
- . 'The Feeling of Knowing: Some Metatheoretical Implications for Consciousness and Control'. *Consciousness and Cognition* 9 (2000): 149–71.
- . 'How Do We Know that We Know? The Accessibility Model of the Feeling of Knowing'. *Psychological Review* 100 (1993): 609–39.
- Lehrer, Keith. *Theory of Knowledge*. Boulder, CO: Westview Press, 1990.
- Lewis, David. 'Elusive Knowledge'. *Australasian Journal of Philosophy* 74 (1996): 549–67.
- . 'New Work for a Theory of Universals'. *Australasian Journal of Philosophy* 61 (1983): 343–77.
- Ludlow, Peter. 'Contextualism and the New Linguistic Turn in Epistemology'. *Contextualism in Philosophy: Knowledge, Meaning and Truth*. Eds. Gerhard Preyer and Georg Peter. New York, NY: Oxford UP, 2005. 11–50.
- Mossler, Daniel G., Robert S. Marvin, and Mark T. Greenberg. 'Conceptual Perspective Taking in 2- to 6-Year-Old Children'. *Developmental Psychology* 12 (1976): 85–6.
- Nickerson, Raymond. 'How We Know – and Sometimes Misjudge – What Others Know: Imputing One's Own Knowledge to Others'. *Psychological Bulletin* 125 (1999): 737–59.
- Nisbett, Richard, Kaipeng Peng, Incheol Choi, and Ara Norenzayan. 'Culture and Systems of Thought: Holistic Versus Analytic Cognition'. *Psychological Review* 108 (2001): 291–310.
- O'Neill, Daniela, Janet Wilde Astington, and John H. Flavell. 'Young Children's Understanding of the Role that Sensory Experiences Play in Knowledge Acquisition'. *Child Development* 63 (1992): 474–90.
- Papfragou, Anna, Peggy Li, Youngon Choi, and Chung-hye Han. 'Evidentiality in Language and Cognition'. *Cognition* 103 (2007): 253–99.
- Plato. *The Theaetetus of Plato*. Trans. M. J. Levett. Ed. Myles Burnyeat. Indianapolis, IN: Hackett Publishing Company, 1990.
- Pust, Joel. 'Against Explanationist Skepticism Regarding Philosophical Intuitions'. *Philosophical Studies* 106 (2001): 227–58.
- Reder, Lynne. 'What Determines Initial Feeling of Knowing? Familiarity with Question Terms, Not with the Answer'. *Journal of Experimental Psychology: Learning, Memory and Cognition* 18 (1992): 435–51.
- Repacholi, Betty and Alison Gopnik. 'Early Reasoning about Desires: Evidence from 14- and 18-Month-Olds'. *Developmental Psychology* 33 (1997): 12–21.
- Ruffman, Ted. 'Do Children Understand the Mind by Means of Simulation or a Theory? Evidence from their Understanding of Inference'. *Mind and Language* 11 (1996): 388–414.
- Satpute, Ajay B. and Matthew D. Lieberman. 'Integrating Automatic and Controlled Processes into Neurocognitive Models of Social Cognition'. *Brain Research* 1079 (2006): 86–97.

- Saxe, Rebecca. 'Against Simulation: The Argument from Error'. *Trends in Cognitive Sciences* 9 (2005): 174–9.
- and Nancy Kanwisher. 'People Thinking about Thinking People: The Role of the Temporo-Parietal Junction in the Theory of Mind'. *Neuroimage* 19 (2003): 1835–42.
- , Laura E. Schulz, and Yuhong V Jiang. 'Reading Minds versus Following Rules: Dissociating Theory of Mind and Executive Control in the Brain'. *Social Neuroscience* 1 (2006): 284–98.
- Schwartz, Bennett. 'Illusory Tip-of-the-Tongue States'. *Memory* 6 (1998): 623–42.
- Shope, Robert K. *The Analysis of Knowing: A Decade of Research*. Princeton, NJ: Princeton UP, 1983.
- Sosa, Ernest. 'A Defense of the Use of Intuitions in Philosophy'. *Stich and His Critics*. Eds. Michael Bishop and Dominic Murphy. Oxford: Blackwell, forthcoming (2008).
- Speas, Margaret. 'Evidentiality, Logophoricity and the Syntactic Representation of Pragmatic Features'. *Lingua* 114 (2004): 255–76.
- Spencer, N. J. 'Differences Between Linguists and Non-Linguists in Intuitions of Grammaticality–Acceptability'. *Journal of Psycholinguistic Research* 2 (1973): 83–98.
- Spicer, Finn. 'Knowledge and the Heuristics of Folk Epistemology'. *New Waves in Epistemology*. Eds. Vincent F. Hendricks and Duncan Pritchard. London: Palgrave Macmillan, 2007.
- Stanley, Jason. *Knowledge and Practical Interests*. New York, NY: Oxford University Press, 2005.
- Vogel, Jonathan. 'Are there Alternatives to the Closure Principle?'. *Doubting: Contemporary Perspectives on Skepticism*. Eds. M. Roth and G. Ross. Dordrecht: Kluwer, 1990. 13–27.
- Weatherson, Brian. 'What Good are Counterexamples?'. *Philosophical Studies* 115 (2003): 1–31.
- Weinberg, Jonathan, Shaun Nichols, and Stephen Stich. 'Normativity and Epistemic Intuitions'. *Philosophical Topics* 29 (2001): 429–60.
- Wellman, Henry M., D. David Cross, and Julianne Watson. 'Meta-Analysis of Theory-of-Mind Development: The Truth about False Belief'. *Child Development* 72 (2001): 655–84.
- , Fuxi Fang, David Liu, Liqi Zhu, and Guoxiong Liu. 'Scaling of Theory-of-Mind Understandings in Chinese Children'. *Psychological Science* 17 (2006): 1075–81.
- Wierzbicka, Anna. *Semantics: Primes and Universals*. New York, NY: Oxford UP, 1996.
- Willett, Thomas. 'A Cross-Linguistic Survey of the Grammaticalization of Evidentiality'. *Studies in Language* 12 (1988): 51–97.
- Williamson, Timothy. 'Contextualism, Subject-Sensitive Invariantism and Knowledge of Knowledge'. *The Philosophical Quarterly* 55 (2005): 213–35.
- . 'Philosophical Intuitions and Scepticism about Judgment'. *Dialectica* 58 (2004): 109–53.
- Wimmer, Heinz and Josef Perner. 'Beliefs about Beliefs: Representation and Constraining Function of Wrong Beliefs in Young Children's Understanding of Deception'. *Cognition* 13 (1983): 103–28.
- Woolley, Jacqueline and Marc Bruell. 'Young Children's Awareness of the Origins of their Mental Representations'. *Developmental Psychology* 32 (1996): 335–46.