

For: *The Routledge Companion to the Philosophy of Psychology*, Paco Calvo, Sarah Robins and John Symons, eds

The Psychology of Epistemic Judgement

Jennifer Nagel and Jessica Wright

1. Introduction: Evaluations of knowledge, belief and justification

Human social intelligence includes a remarkable power to evaluate what people know and believe, and to assess the quality of well- or ill-formed beliefs. Epistemic evaluations emerge in a great variety of contexts, from moments of deliberate private reflection on tough theoretical questions, to casual social observations about what other people know and think. We seem to be able to draw systematic lines between knowledge and mere belief, to distinguish justified and unjustified beliefs, and to recognize some beliefs as delusional or irrational. This article outlines the main types of epistemic evaluations, and examines how our capacities to perform these evaluations develop, how they function at maturity, and how they are deployed in the vital task of sorting out when to believe what others say.

1.1 Knowledge and belief attribution

At a formal dinner party, your hand slips and you knock over a full glass of red wine. Silence falls. Your sense of awkwardness is immediate, and very different in character from the feeling you'd get from a similar moment of clumsiness while home alone: your embarrassment in the social setting is made possible in part by your instant capacity to register that others know that you were the one who did this. Human beings are naturally alert not only to what others know, but also to what they think: in a variation of our scenario, imagine a situation in which the other guests are distracted at the moment when your neighbour hits your wine glass, and then they all look over at you, wrongly believing that you are the culprit behind the spreading stain on the tablecloth. Here again you'll have an instant emotional reaction as you register their misunderstanding. But how could you detect so swiftly what your dining companions knew or wrongly thought to be the case, even before any of them had the chance to speak? How could you read their minds?

By reviewing some basic structural features of the states of knowing and of thinking that something is the case, we can start to make sense of the human capacity to register these states. Epistemologists disagree about the finer details, but certain core features of the relationship between knowing and believing are widely recognized. It is widely agreed that both knowledge and belief are states which link an agent to a proposition (the verbs “know” and “believe” have other senses as well, but here we’ll focus on propositional knowledge and belief). Read broadly, the term “agent” includes anything that can act, including groups and non-human animals; for simplicity in what follows, our main examples of knowing and believing agents will be individual human beings. Propositions are commonly understood to be packages of information, packages like the meanings of natural language sentences, which serve as “the primary bearers of truth and falsity” (King, Soames, & Speaks, 2014, p.5). Propositions can be expressed directly in natural language, using “that-” clauses (e.g. “the host thinks that I broke the wine glass”; “my friend knows that I am seated on his left”). They can also be expressed indirectly, as embedded questions or “wh-” clauses: (“the host knows what was in the glass”, “my friend knows where the glass fell”). The proposition that is known here is the true answer to the embedded question (“there was Pinot Noir in the glass”, “the glass fell beside his plate”), whatever that answer might be (Karttunen 1977). Our ability to use embedded questions to express a proposition indirectly enables us to characterize others as knowing things that we don’t know ourselves. For example, you can say or judge, “this tall man knows what is on the far side of this wall”, where the proposition “the goat is on the far side of this wall” is the thing he knows and you don’t. We can keep track of whether others share our knowledge, or know more or less than we do; we can also keep track of what they believe, whether it is true or false.

One major point of contrast between knowledge and belief concerns the range of truth values available for believed and known propositions: philosophers since Plato (369BCE/1990) have accepted that what is believed can be true or false, but what is known can only be true (or factual). The other dinner guests can believe that you knocked over the wine glass when you didn’t, but they cannot know something that was

not in fact the case. Of course, we sometimes mistakenly take someone to know something which later turns out to be false, but as soon as we discover the falsity we have to retract the judgment that the agent knew it to be the case.

Knowledge attributions are expressed in language not only with the use of the verb “to know”, but also with a broader family of related verbs. Attitude verbs like “to know” which can take only true complements are called “factive” verbs: other factives include “realize that”, “notice that”, “see that” and “be aware that” (Kiparsky & Kiparsky 1970). Non-factives like “believe that”, “think that”, “hope that” and “doubt that” can take both true and false complement clauses. The line between factives and non-factives seems to be drawn in all languages, with significant characteristics of this distinction appearing cross-linguistically (Egre 2008; Lahiri 2002).

Extending the terminology of linguistics into epistemology, attitudes that can be held only to truths are called factive attitudes; some of these attitudes (like knowing) constitute states, where others (like learning or forgetting) constitute events or processes. According to one influential account of knowledge, knowledge is the most general factive mental state: on this account, realizing, noticing, perceiving, being aware or seeing that p is the case all entail knowing that p (Williamson 2000). There is some controversy about this account of knowledge: some deny that knowledge is best characterized as a mental state (e.g. Fricker, 2009), and others raise doubts about whether *all* factives entail knowing, particularly emotive factives such as “is happy that” (Fantl 2015). But it is broadly agreed that the verb “to know” is not the only verb denoting the possession of knowledge, and that factives like “is aware that” and “recognizes that” also function to attribute knowledge. It is also widely agreed that it is essential to the state of knowledge that the knower is getting things right.

It is clearly possible to believe without knowing, notably in cases where one’s belief is false. But is it possible to know without believing? Most contemporary epistemologists take knowledge to entail belief: if

Jones knows that the meeting is cancelled, then it follows automatically that Jones believes that the meeting is cancelled (Steup & Ichikawa 2008). This entailment thesis is sometimes challenged. Colin Radford (1966) introduced an intuitive example involving a hesitant student being quizzed on historical facts. When asked for the date of Queen Elizabeth the First's death, the student very hesitantly answers "1603." The student reports that he feels he is just guessing, and has no clear sense of which of his answers are likely to be right, but produces a reasonable number of accurate answers, which Radford argues he must be remembering, and therefore must know. On the basis of several such examples, Radford ultimately concludes that knowledge requires neither confidence nor even belief. Radford's examples did not originally win a warm reception: philosophers have suggested that these cases do actually manifest belief, or that they fail to manifest knowledge, or have argued that they are equally dubious or borderline examples of both knowledge and belief (e.g. Armstrong, 1969; Lehrer, 1968).

Radford's challenge has been revived in recent years, as experimental philosophers have used more systematic methods to test attributions of knowledge and belief. In a between-subjects design, Blake Myers-Shulz and Eric Schwitzgebel (2013) probed responses to a case directly modelled on Radford's, and to other cases involving ambivalent, distracted and self-deceived agents, generally agents who are now failing to act on the basis of what they have at some point learned. Participants read these cases, and were asked either a question of the form "Does [name] know that [key proposition]?" or "Does [name] believe/think that [key proposition]?" In all of the five cases they tested, the group asked about knowledge was more willing to attribute it than the group asked about belief. The between-subjects design made it hard to measure compliance to the thesis that knowledge entails belief, but the authors extrapolated that only about half of their participants made judgments in line with that principle for these cases.

There are several ways to respond to these results: one possible response is to insist that the state of knowledge actually does entail belief, while seeing experimental participants' intuitive responses as compromised or distorted in these interesting and complex cases. Perhaps participants are more cautious

or thoughtful when asked what ambivalent characters believe or think, than when asked what they know, and these borderline cases show something about our variable epistemic sensitivities rather than about the nature of knowledge and belief. It is also possible to argue that lay epistemic judgments do generally reflect the principle that knowledge entails belief, and instead criticize the methods that have produced apparently contrary evidence. Perhaps there are several senses of “believe”, and if we direct attention to the right sense of the word, we’ll see results more in line with the traditional view that knowing entails believing. One hint that there may be some ambiguity in “believe” which is activated in these cases is that Myers-Schulz and Schwitzgebel got markedly different responses to some of their cases when they used the verb “think” as opposed to “believe”. In the sense which takes a propositional complement, these verbs are typically interchangeable in English: “John thinks that the ball is in the box” and “John believes that the ball is in the box” usually register as equally true (or false). But especially in contexts where people are feeling torn about what to say or do, the verb “believe” may have more vivid connotations of what is currently in mind. For Myers-Schulz and Schwitzgebel’s case involving a self-deceived husband whose wife was cheating on him, and for their more dramatic version of Radford’s student case, in which the student (Kate) was extremely stressed under time pressure, and found her mind blank, participants were especially reluctant to describe the agent as believing that the key proposition was true, although a clear majority were willing to describe the agent as thinking it was so, as well as knowing it was so.

According to David Rose and Jonathan Schaffer (2013), the sense of “believe” that matters for epistemology is a dispositional rather than an occurrent (activated) sense: just as you can have standing knowledge on a point not currently in mind, you can have a standing belief which isn’t conscious at this moment. They retested the Myers-Schulz cases with questions designed to elicit that sense of “believe”, for example, “Did Kate still believe (in the sense that she still held the information in her mind even if she could not access it) that Queen Elizabeth died in 1603?” When prompted in this manner, most participants reverted to complying with the principle that knowledge entails belief. In a somewhat similar spirit, Wesley Buckwalter and colleagues (Buckwalter et al. 2015) distinguished a thin and a thick sense of

“believe”, where the thin sense only requires the storage of information, and the thick sense involves active endorsement and the control of action. Generally, these two senses are aligned, but in special cases, like the ambivalent cases probed by Myers-Schulz, they can come apart. Buckwalter and colleagues elicited attributions of “thin” belief by asking questions of the form “At least on some level, does [agent] think that [key proposition]?” Similar probes tested for thin knowledge attributions. Under these conditions, an overwhelming majority of participants showed intuitions complying to the principle that knowledge entails belief.

To accommodate those who worry that delicate scenarios of ambivalence and self-deception may constitute exceptions to the general rule that knowledge entails belief, we’ll set those aside and focus on simpler core cases, where what is known is also believed. We have observed that what is known must be true, but not every true belief constitutes knowledge: lucky guesses and wishful thinking can be true while falling short of knowledge. Knowledge requires some sort of ‘good basis’, and cannot be arbitrary and unjustified. It remains controversial in epistemology exactly what kind of justification or basis is required for knowledge, but we can start by examining our natural resources for making basic judgments about evidence, justification and the quality of judgments.

1.2 Justification and evidence collection

Some judgments are careful, patient and well-considered; others are hasty and confused. Epistemologists’ stock examples of justified beliefs include perceptual judgments in favourable circumstances, sound inferences and responsible learning from testimony; stock examples of unjustified belief include wishful thinking and hasty generalization. There is evidence that laypeople agree with epistemologists about stock examples of justified and unjustified beliefs, and also agree that some false judgments can nevertheless be justified, for example where agents judge carefully on the basis of extensive but misleading evidence (Nagel et al. 2013). There are various theories of how we evaluate the judgments of others as justified or unjustified. One prominent theory suggests that we do so by comparing these judgments to learned

stereotypes of reliable and unreliable belief formation (Goldman 1994); others have argued that we ultimately evaluate not in terms of reliability, but in terms of some other factor, such as ability (Turri 2016).

Justification comes in degrees, and the question what constitutes a good enough basis for a belief to count as adequately justified is a difficult one. Rival theories are generally agreed that having more evidence for a proposition tends to improve one's level of justification, but how much evidence does one need to collect before one is justified enough to count as knowing? If you are wondering whether your bank will be open this Saturday, is it enough to remember a Saturday visit at some point in the past, or do you have to double-check the website, or visit the branch to make sure? We have some reason to expect variation here, depending on the circumstances of the person making a judgment. Work on adaptive decision making has shown that when stakes are high we instinctively collect more information before making up our minds, and we also expect others to do the same (see Nagel, 2010, for a review). We also adapt our methods of inquiry, switching from more automatic, intuitive ways of making judgments to more reflective and deliberate ways. We search for less information in easy evidential climates, where we learn that we can get the right answer by looking at a few cues, and we search much more extensively in climates where we discover that initial impressions can be misleading (Lee et al. 2014). If we are evaluating the epistemic states of someone in a high-stakes situation, we will naturally expect them to be more cautious, and to collect more evidence than their low-stakes counterpart, before arriving at a firm belief. Given our current working assumption that belief is one of the necessary conditions for knowledge, we will expect high-stakes thinkers to work harder than their low-stakes counterparts to achieve a state that we count as knowledge.

1.3 The social neuroscience of epistemic judgments

Judging whether someone knows or believes something can involve a great range of cognitive functions: we take in various perceptual cues, we calculate the visual perspectives of others, we process language,

either in a story about the agent or in listening to what the agent is saying, we draw on background information about the world in general and social scripts in particular, and we make various types of explicit and implicit inferences (for a review, see Apperly, 2011). With such a diverse array of mental functions feeding into the attribution of knowledge and belief, there is nevertheless a surprisingly high level of specialization in the brain for the attribution of these epistemic states.

Functional neuroimaging has extracted an increasingly sharp picture of the brain regions implicated in epistemic evaluations, by matching epistemic judgment tasks against control tasks closely matched for their demands on language, perspective taking and other dimensions of appraisal. Within the range of brain regions involved in social cognition (including perception of agency and emotion), several regions seem especially significant for epistemic evaluation, most notably the medial prefrontal cortex (mPFC) and the bilateral temporo-parietal junction (TPJ). The mPFC is thought to be activated in a wide range of social cognition, including thoughts about human appearances, traits, sensations and personalities (Saxe & Powell 2006); it also seems to play a significant role in keeping track of the reliability or ‘social value’ of others, for example in playing strategic games with partners over time (Behrens et al. 2008). While the mPFC was once thought to serve as the ‘decoupling mechanism’ for separating representations from reality in our evaluations of others, enabling the representation of misconceptions (Frith & Frith 2003), there is increasing evidence that this function is handled instead by the TPJ, together with reasoning about knowledge or true belief (Aichhorn et al. 2009). Across fMRI studies, the TPJ is consistently activated in reasoning about the knowledge and beliefs of others (Schurz et al. 2014); sharp deficits in epistemic state reasoning are produced when this region is temporarily impaired through transcranial magnetic stimulation (Young et al. 2010), or more permanently impaired following a stroke or injury (Apperly et al. 2004).

Researchers probing the nature of epistemic evaluations have a number of tools at their disposal. For example, we can make epistemic evaluations either in the course of verbal tasks such as reading

hypothetical stories about agents, or in non-verbal tasks such as watching silent films, or playing strategic games with live partners. Interestingly, it does not seem to matter whether we are cued by linguistic or non-linguistic stimuli: a comprehensive meta-analysis of imagining studies found no significant differences in activation between verbal and non-verbal Theory of Mind tasks (Carrington & Bailey 2009). This result matters to epistemologists who may be wondering whether evaluations of narrative thought experiments—like the Gettier case scenarios discussed in section 3—are managed differently from epistemic evaluations in the wild. Physiologically, these evaluations seem to work at least roughly the same way.

There is nevertheless a complex relationship between language and epistemic evaluation. Children's linguistic abilities are highly correlated with their ability to pass tests of explicit false belief attribution, with early linguistic ability an even stronger predictor of early mindreading ability than the converse (Milligan et al. 2007). The causal dependencies here are hard to calculate, however, not least because explicit false belief tests place significant demands on linguistic competence, including managing the odd pragmatics of being asked a question about belief by someone who already knows the answer (Helming et al. 2014). Interestingly, adults with severely impaired syntax are able to perform very well on first- and second-order mental state attribution tasks presented pictorially (Apperly et al. 2006; Varley et al. 2001). Adults can make intelligent judgments about others' knowledge, beliefs, and even others' false beliefs about the false beliefs of third parties, despite an inability to understand spoken that-clauses, or to distinguish between the sentences "the fish was eaten by the cat" and "the fish ate the cat". Even if language plays a key role in the development of our capacity for epistemic evaluations, and even if adults in the non-clinical population often use language to express their epistemic evaluations, competence in epistemic evaluation is interestingly independent of competence in grasping the syntactic structure of language.

Some epistemic evaluations are executed deliberately: for example, members of a jury may be instructed to reflect on the question of whether the defendant suspected or knew that the gun was loaded. Epistemic

evaluations are also performed without specific instruction, and there is a lively debate about the extent to which we make such evaluations automatically or spontaneously. Are we invariably and involuntarily triggered to track what others around us see, know, and believe? Researchers have probed this question in a number of ways. In the dot perspective task, participants need to count the number of dots presented in a virtual room on screen in front of them: sometimes they see an image of a person standing in middle of the virtual room looking at the dots on one side of the room, with his back to either more dots or a blank wall behind him. Participants are asked how many dots they themselves can see, and respond more swiftly if the on-screen avatar can see that same number of dots; when the image of a person is replaced with an image of a similarly-sized column, the effect vanishes (Samson et al. 2010). This and various other effects of interference from the perspectives of others have led some researchers to believe that we automatically devote mental resources to computing what others can see, even when this is not relevant to our current task. Not all researchers are persuaded: some have argued that the effect has more to do with domain-general direction of attention, discovering that the “perspective interference” effect appears again if we replace the virtual person with a directional arrow rather than a column (Santesteban et al. 2014). However, there is continuing controversy over whether the underlying basis of our response to what others can see is different in extent or kind, when compared with well-matched control conditions (Catmur et al. 2016; Schurz et al. 2015). In any event, whether or not we have a special domain-specific response to the minds of others that is invariably triggered by witnessing other agents, it is widely agreed that epistemic evaluations of one kind or another are performed very frequently, and rapidly, in ordinary social navigation and conversational exchanges.

2. The development of knowledge and belief attribution

2.1 Developing a capacity for explicit epistemic evaluations

When and how do human beings gain the ability to evaluate epistemic states? To answer this question, researchers have studied children’s competence in answering explicit questions about knowledge and

belief, and questions about how agents with different epistemic positions will act; they have also looked at spontaneous use of mental state terms, and at subtler patterns of anticipatory looking and surprise, in reaction to agents who are seen to either witness or fail to witness changes in the world that are relevant to the actions they are about to perform.

One of the most clearly established findings concerns the emergence of an ability to pass the explicit false belief test, usually between ages four and five (Wimmer & Perner 1983). In this test, the child sees an agent placing an object (such as a bar of chocolate) in a certain location (say, a cupboard); the agent then briefly leaves the scene. During her absence, the chocolate is moved (say, to a drawer) as the child watches. The child is asked where the agent will look for the chocolate when she returns. Around the time that they turn three, children typically fail this task: a large-scale meta-analysis found roughly two-thirds of 36-month-olds saying that the agent will look in the closed drawer, where the chocolate is now located; by 44 months about 50% are answering correctly, expecting the agent to look in the cupboard where she last saw the chocolate, and a year later about 75% produce the right answer (Wellman et al. 2001). Similar progressions are seen in children from different cultures, and appear whether the task is presented with live actors, puppets or videos.

One way of understanding the switch from failing to passing the false belief test is to say that it is only in the fourth year of life that children develop the concept of belief (Rakoczy 2015; Perner 1991). Before this age, children can see agents as having some mental states—notably goals, such as the goal of getting the chocolate—but they seem to be unable to represent agents as having states of mind which may be out of line with reality. Children are sometimes described as having a natural bias towards egocentrism in their epistemic evaluations: they find it hard not to project their own knowledge onto others, a tendency which also appears, in subtler ways, in the epistemic evaluations of adults (Birch & Bloom 2004; Keysar et al. 2003).

Interestingly, the capacity to attribute knowledge, a state which is always in line with reality, seems to emerge before the capacity to attribute belief. Children can explicitly attribute knowledge and ignorance by the age of three (Pratt & Bryant 1990). Across cultures, children use words meaning “know” earlier and more often than they use words meaning “believe” (Tardif & Wellman 2000; Shatz et al. 1983; Pascual et al. 2008). Some researchers had been inclined to dismiss children’s very early talk of knowledge—especially their use of “I don’t know” (or “I dunno”)—as merely formulaic, but more recent cross-cultural research has supported the idea that by the age of two, children speak of knowledge in a way that is sensitive to what is known, and show similar patterns in describing their own lack of knowledge and in asking others what they know (Harris et al. 2016). Describing the knowledge that agents possess or lack is in one important way simpler than describing the beliefs that they have: the natural domain of what can be known is reality, where the natural domain of what is believed is much wider, because people can believe virtually any imaginable truth or falsehood. If knowledge is a window onto reality (which may be blocked in cases of ignorance), belief is a depiction of something which may be reality-independent; in order to track the false beliefs of others, one must devote some cognitive resources to representing those belief contents, decoupled from reality and tied to the relevant agents, alongside a rival representation of reality itself. As adults we find it easy to notice and report the misconceptions of others; by reviewing the difficulties that young children seem to have on this frontier, we can come to appreciate how difficult it is to perform explicit epistemic evaluations of mistaken agents. At present, the capacity to track false beliefs appears to be uniquely human: while non-human primates show some capacity to track whether other members of their species have knowledge, they have not shown any ability to represent false belief (Kaminski et al. 2008; Marticorena et al. 2011).

2.2 Implicit epistemic evaluations

Well before they can give accurate explicit answers to questions about knowledge and belief, infants and children show subtler signs of tracking epistemic states. Three-year-olds who give the wrong answer to the explicit false belief test nevertheless consistently look at the right location as they anticipate the returning

agent's search for the chocolate (Clements & Perner 1994). Researchers have suggested that even younger children show signs of implicit belief attribution: in a "violation of expectancy" paradigm, 15-month-olds are reported to look longer, as if in surprise, at returning agents who search for an object at its new and actual location rather than at the last place the agent saw it (Onishi & Baillargeon 2005). Implicit false belief attribution is gauged not only by means of anticipatory looking direction and looking time, but also in helping tasks, where it seems 18-month old toddlers can modify their behavior to help agents who may have false beliefs (Buttelmann et al. 2009). Findings for knowledge are even more striking: looking time measures have also been taken to show that infants seem to be able to track the difference between other agents' knowledge and ignorance as early as 12.5 months (Luo & Baillargeon 2007).

Various rival theories attempt to explain the early infant findings and the reasons why explicit false belief attribution lags apparently early sensitivity to false belief. Some have argued that the infant results should not be understood as signs of genuine epistemic evaluation, but as low-level responses to perceptual patterns in the stimuli (Heyes 2014) or as the products of domain-general statistical learning processes, guided by infant biases to attend to human faces and actions (Ruffman 2014). Others argue that infants really do see others as knowing and believing, possessing a repertoire of innate or early-developing mental state concepts whose expression is initially blocked in explicit tasks for reasons having to do with factors such as processing load or the difficult pragmatics of responding to questions about false beliefs (Carruthers 2016; Helming et al. 2014). Some see children as progressing from a stage at which they use a more basic subsystem to attribute goals and knowledge, to a later stage, at which a distinct subsystem makes belief attribution possible, first implicitly and then explicitly (Baillargeon et al. 2010). Others have a more radical two-systems view, in which infants initially track something more basic than full-blown propositional attitudes: perhaps infants first encode just relationships between agents, objects and locations (Butterfill & Apperly 2013), or records of the objects and events associated with the agent (Perner & Roessler 2012), and only much later become capable of the flexible and inferentially rich theory of mind

which propositional attitude attribution makes possible. Advocates of the two-system view argue that the more basic infant system is retained in adults, and used for rapid judgments in daily social navigation.

Whether genuine epistemic evaluations are seen as early or late-developing, all researchers face the challenge of explaining how it is that we come to be able to compute the contents of the knowledge and beliefs of other agents. There is a general consensus that our grasp of these contents begins with our ability to detect others' perceptual access to reality, but there is ongoing controversy about how we move from seeing others as knowing about certain aspects of reality, to seeing them as having potentially false beliefs about various other ways the world might be. Ongoing research into the emergence of our capacities to attribute knowledge and belief promises to shed further light on the nature of these capacities and the relationship between these epistemic states themselves.

3. Mature epistemic evaluations: biases and paradoxical patterns

3.1 Shifting intuitions about knowledge

Some patterns of instinctive epistemic evaluation are surprising, counter-intuitive, even paradoxical. One of the most striking such patterns is associated with skepticism. We are ordinarily quite liberal in our positive attributions of knowledge: “know” is one of the ten most common verbs in English (Davies & Gardner 2010), and a quick corpus search of spontaneous use of this verb in conversation and writing shows that it most commonly appears without negation. Reading a story about an ordinary man in a furniture store looking at a bright red table, and asking his wife whether she likes this red table, you wouldn't ordinarily hesitate to say that he knows that the table is red; indeed, this looks like a paradigm case of knowing. Still, there are ways of getting people to hesitate to attribute knowledge here. You could, for example, mention that a white table under a bright red spotlight would look exactly the same to this shopper, and he hasn't checked whether the lighting is normal. Even if you make it clear to the reader that the lighting is normal and the table is red, knowledge attribution rates plummet as that problematic

hypothetical possibility is mentioned, with only a minority still seeing the shopper as knowing (Nagel et al. 2013).

Virtually any case of ordinary empirical knowledge can be re-written as a “skeptical pressure” case by mentioning some problematic possibility. Does a woman who glances at a working clock know what time it is? Of course. But if we emphasize that although this particular clock is working, sometimes clocks are broken, and that our woman doesn’t look at the clock for long enough to be sure that it is operating correctly, there is considerable reluctance to see her as really knowing what time it is on the basis of her glance at the clock. What is interesting about the shifts in these cases is that they simply highlight something we in some sense knew all along (such as the fact that clocks are sometimes broken). When attention is directed to this fact, the case no longer registers as a very good intuitive case of knowledge.

By describing a situation in slightly different ways, we can manipulate whether it seems like a case of knowledge or ignorance. We can also shift ourselves from thinking to doubting that we know anything at all about our environment by focusing on problematic possibilities concerning large-scale deceptive scenarios (could I be a brain in a vat?). Some philosophers have proposed subtle theories about the semantics of “knows” to explain these shifts; others have argued that knowledge itself is more complex than it might appear, or that something is going wrong in our intuitive impressions of knowledge in these cases, just as something goes wrong in our intuitive impressions of the length of lines in the Müller-Lyer illusion.

The dominant semantic theory is known as contextualism. According to contextualists, “know(s)” is a context-sensitive term, like “here”, “tomorrow” or “tall” (DeRose 2009). There is no fixed height that makes something count as tall: what it takes to be tall for a rose bush is rather different from what it takes to be tall for a redwood. Implicit comparison classes seem to be called into play, with “tall” always meaning something like “relatively great in vertical magnitude compared to things of its class.” The same

thing can be judged different ways in comparison to two classes of which it is a member: a man can be tall for an American but not tall for a basketball player. Applied in the epistemic realm, contextualists say that we have variable epistemic standards: someone can count as knowing by the standards of the casual bus-stop conversation, without counting as knowing by the standards prevailing in a court of law (Stine 1976). There are various theories of how these shifting standards function: some argue that what shifts is how much evidence we need (Cohen 1999), while others argue that what shifts is how large a range of possibilities of error we need to eliminate (Lewis 1996). However standards are operationalized, contextualists typically hold that the average John Doe can be speaking truly when he says, in a causal context, “I know that there is beer in my fridge”, while the skeptic who looks down on John Doe from a discussion about brains in vats can also be speaking truly when he says, “John Doe does not know that there is beer in his fridge.” These two speakers are not contradicting one another, according to contextualism, because they are not latching onto the same thing with their use of the word “know”.

Contextualism faces a number of challenges. For ordinary context-sensitive terms like “tomorrow”, there is a relatively simple and easily-learned recipe for calculating the content expressed by the term in any given context of use. If “knows” is context-sensitive, it is not so obvious how this sensitivity works; critics of contextualism have argued that it does not exactly fit the pattern of gradable adjectives like “tall” (Stanley 2005), while defenders have suggested that “know(s)” must have its own distinctive type of context-sensitivity (DeRose 2009). To date, contextualists have provided only partial accounts of the circumstances under which epistemic standards shift one way or another, and the search for alternate explanations continues.

One rival explanation of shifts in knowledge attribution focuses not on the attributor of knowledge but on the knowing subject. Some of the most discussed cases of apparently shifting standards for knowledge concern shifts from low- to high-stakes situations. For example, a person with casual curiosity about whether a flight will have a layover in Chicago could seem to know that it will stop there after a quick

glance at a schedule, where a person for whom the location of the layover is a life-and-death matter would need to do more research in order to come across as knowing (Cohen 1999). According to the position now known as “Interest-Relative Invariantism” (IRI), the verb “know(s)” always denotes the same relation between an agent and a proposition, but it is a relation in which stakes matter directly as a factor in whether someone has knowledge (Stanley 2005). Advocates of this view have aimed to support it by appealing to patterns of intuitive epistemic evaluation, and also by appeal to more abstract principles about the relationship between knowledge and action. Meanwhile, critics of IRI have offered alternative accounts of the relationship between stakes and knowledge, for example, focusing on the impact that stakes have on confidence (Weatherson 2005; Bach 2005; Nagel 2010a). Critics of IRI have also noted that the theory does not seem to explain superficially similar shifts in apparent knowledge, like those occurring in the clock case, where nothing much needs to be at stake in order to produce the skeptical impression that the person glancing at the clock fails to have knowledge.

Another strategy to handle conflicting or paradoxical patterns of epistemic intuition is to argue that some intuitions are the products of illusion or bias. Some philosophers have argued that rising stakes switch our attention from ordinary ground-level questions of what is known up to higher-level questions concerning iterations of knowledge (do I know that I know this?), distracting us from what is really at issue (Williamson 2005). Others have argued that by mentioning hypothetical possibilities of error we bias the focus of epistemic evaluation in a distorting fashion, in one way or another (Gerken 2013; Nagel 2010b). Shifting patterns of epistemic evaluation remain an open source of controversy in epistemology, philosophy of language, and philosophy of mind.

3.2 Gettier cases

The person who looks at a broken clock doesn’t often come away knowing what time it is. But imagine someone who looks at a stopped clock at just the right moment to form a true belief. Does she know, as she looks at the hands of the broken clock pointing to 3:15, that it is now 3:15? Notice that this is not only

a true belief, but a justified true belief, in some sense of “justified”. The person who looks at a clock in order to tell the time is behaving responsibly and collecting evidence relevant to the question she is settling (and indeed in luckier circumstances, with a functioning clock, we’d be happy to count her as knowing the time).

The case of the accurate broken clock is generally seen as a case of justified true belief without knowledge, a type of scenario known in the literature as a “Gettier case”, after Edmund Gettier, who included several such cases in a short paper arguing against the idea that knowledge could be analyzed as justified true belief (Gettier 1963). Gettier’s counterexamples spurred the development of a series of more elaborate theories of knowledge, in which additional conditions were added to justified true belief, as well as fresh analyses of knowledge in terms of causal and counter-factual conditions (Shope 1983). To date no proposed analysis of knowledge has won very widespread acceptance among epistemologists, and some now argue that it is unlikely to be possible to analyze knowledge into true belief plus non-epistemic factors (Williamson 2000).

Whether they are trying to forge a new analysis of knowledge, or trying to show that no such analysis will be satisfactory, epistemologists often rely on intuitive assessments of particular cases like the stopped clock case as evidence about the nature of knowledge. Concerns are sometimes raised about these subtle epistemic evaluations. In one heavily-cited paper from 2001, three critics of intuition-driven philosophy argued that epistemic evaluations of Gettier cases might show more about local cultural practices than about the nature of knowledge itself (Weinberg et al. 2001). On the basis of a small-scale study of American undergraduates of different ethnic backgrounds, they concluded that Gettier case evaluations differed by cultural group, with East and South Asian participants attributing knowledge where Western participants saw here true belief. These conclusions were based on responses to the following case:

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?

REALLY KNOWS

ONLY BELIEVES

Subsequent researchers failed to replicate the original findings of cultural variation for this case, finding instead that most respondents across different cultural backgrounds evaluated it as a case of mere belief, in line with the traditional philosophical view that cases like these show justified true belief without knowledge (Kim & Yuan 2015; Seyedsayamdost 2015; Nagel et al. 2013). A much larger study of participants in Brazil, India, Japan and the USA found similar responses to a variety of Gettier cases across all of these cultures, prompting the conclusion that these responses are part of an innate and universal “core folk epistemology” (Machery et al. 2015). The extent to which intuitions about knowledge are shared across cultures, and the ultimate basis of these intuitions, are questions for further research.

4. Epistemic Vigilance

Much of what we learn from the world comes from what others tell us. Hand in hand with this fact, however, comes a hard truth—for a variety of reasons, what others tell us is not always going to be true. In response to this social reality, we have various mechanisms for “epistemic vigilance”, or sorting good from bad information (Heintz et al. 2010). Research on epistemic vigilance has both descriptive and normative dimensions. The descriptive side concerns what cognitive mechanisms, if any, we use for epistemic vigilance. The normative side concerns how these mechanisms are evaluable by epistemic standards, like knowledge; and moral standards, like justice. Both descriptive and normative issues arise on two different fronts: vigilance towards the *informant*, and vigilance towards the *information* communicated.

4.1. Vigilance Toward the Informant

Imagine that you are at an art museum and a very young child approaches you and tells you, with great sincerity, that the museum has made a mistake in taking Picasso to have painted the work of art you are looking at. If your reaction is to express surprise or doubt in response to this statement because of her age—as it likely would be—this is an instance of you exhibiting vigilance towards the *source*.

Vigilance towards the source of testimony is perhaps the most obvious type of epistemic vigilance.

Vigilance towards the source can manifest itself in a number of different ways. We can doubt that the person *knows* or is competent to know (even if she thinks she knows), we can doubt that the person is being sincere, and we can judge the informant on the basis of generalized heuristics, like a sense of trustworthiness. One of the ways that we manifest this tendency is in thinking of an informant as *trustworthy* or *untrustworthy*, *credible* or *not credible*. Such judgments can be made very quickly. In a study probing the evaluation and attribution of traits like trustworthiness and competence, researchers found that after a glimpsing a face for just one tenth of a second, participants were able to form judgments of trustworthiness that correlated very closely with judgments of trustworthiness made without time pressure, more closely than corresponding judgments of attractiveness (Willis and Todorov 2006). Young children already exhibit source-directed vigilance. In particular, like adults, young children see informants as having traits such as trustworthiness, competence and accuracy (Cogsdill et al. 2014; Kuhlmeier et al. 2003).

In attributions of trustworthiness or credibility as a trait to the source, one is judging not the quality of the *information* transmitted, but rather the quality of the *informant*. This move of attributing a trait to an informant may be problematic, if it is an instance of the “fundamental attribution error” (Ross 1977), where agents downplay situational factors and overestimate the contribution of personality and dispositional traits to an agent’s behaviour (O’Sullivan 2003; Gilbert & Malone 1995). Even more problematically, judgments of trustworthiness, competence, and credibility may be based on things that are not epistemically relevant, like the shape of a face (Zebrowitz et al. 1996).

When it is not too costly to do so, agents will exert other kinds of source-centered vigilance. Vigilance towards the source requires that we be aware not only of the agent's explicit interests and motivations, but also who she is in relation to the information that is being communicated to us, and in relation to us. In particular, vigilant agents need to know if the source is competent (i.e., does she know or have reason to believe the information?), and if the source is warm or benevolent (i.e., is she well-intentioned towards you?). Across cultures, evaluations of competence and warmth play a leading role in the interpretation of social behavior (Fiske et al. 2006).

Whatever situation-based or trait-based mechanisms we use in ordinary social navigation, they are hard-pressed to detect deliberate deception. Studies have shown that people, even those trained to detect lies, are remarkably bad at telling when someone is lying (Ekman & O'Sullivan 1991; Bond & DePaulo 2006). The verbal and non-verbal cues most most often read as markers of lying (such as avoidance of eye contact), are not generally valid signs of deception (Sporer & Schwandt 2007; Sporer & Schwandt 2006).

These considerations together paint a dim picture of source-directed epistemic vigilance. However, research on our limitations needs to be balanced against research on our strengths in source-centered vigilance. The more successful side of our basic capacities for epistemic vigilance has been explored most extensively in research on children's learning. It was once thought that children under a certain age are naive epistemic agents—that they will believe anything an informant tells them. Recent work on the topic, however, shows that children are surprisingly discerning in their epistemic evaluations (Woolley & Ghossainy 2013; Vanderbilt et al. 2014; Clément 2010). By the age of two, children prefer to learn from informants who were previously accurate (Clément et al. 2004; Koenig & Harris 2005; Koenig & Woodward 2010), and by four, children will be vigilant towards informants on the basis of probabilistic evidence (Pasquini et al. 2007). Young children will also take into account how knowledgeable their source is (Sabbagh & Baldwin 2001; Sobel & Kushnir 2013; Tummeltshammer et al. 2014; Stephens et al.

2015). Many psychologists now think that children, rather than being gullible, are in fact quite intelligently vigilant towards sources of testimony.

4.2 Vigilance Towards the Information: Coherence

Agents also exhibit epistemic vigilance in relation to the information transmitted to them. The main criterion for vigilance relative to the information itself is whether the information is *coherent* with one's other beliefs. Developmental psychologists have found sophisticated coherence-checking behaviors in children, looking at their responses to labeling errors, inaccuracy, grammatical errors, and inconsistent, illogical and improbable statements (Stephens et al. 2015). 8-month-old infants already prefer to follow the gaze of a reliable looker—one who consistently looks at objects on a screen, as opposed to someone who inconsistently does so (Chow et al. 2008; Tummeltshammer et al. 2014). This is thought to show that infants track coherence (Stephens et al. 2015). Adults also track the coherence of what they are being told, perhaps in part by means of an intuitive sense of how fluently new information can be processed (Reber & Unkelbach 2010).

4.3 Epistemic Injustice

After examining the mechanisms behind epistemic vigilance, we can wonder whether these social mechanisms are operating in the way that they ought to. For example, mechanisms for social vigilance are open to socially learned generalizations and associations that may compromise our epistemic evaluations, either epistemically or morally, in various forms of “epistemic injustice” (Fricker 2007).

Gender is one representative area of concern. Studies on the credibility of expert witnesses are potentially illuminating here. In courtroom decisions, juries may be swayed by gender stereotypes in assessing the credibility of male and female expert witnesses. For example, one team of researchers studied responses to transcripts of a price-fixing case set either in a domain that was either traditionally male (automotive services) or female (cosmetics), with key testimony in otherwise identical cases ascribed to male and female

expert witnesses; they found higher evaluations of the expert witnesses who were testifying in their gender stereotype-congruent domains, and in the traditionally female domain, higher damages awarded to the side represented by the female expert witness (McKimmie et al. 2004). These studies do not show a clear-cut ranking in the perceived credibility of female and male expert witnesses; indeed, other studies have found results cutting against the stereotype-congruity findings (Couch & Sigler 2002). The credibility of male and female expert witnesses depends on a number of different situational factors, including the domain of the case, the complexity of the testimony, and the extent of deliberation (Neal 2014).

If an expert's actual knowledge or objective credibility on a subject has very little to do with his or her gender and the congruency of his or her gender with her field, then instinctive mechanisms for epistemic vigilance can lead us to make epistemically defective or morally problematic judgments about our informants. The kinds of problematic generalizations that matter are ones like 'car mechanics tend to be men'. When faced with a woman who is testifying about car repair issues, we may dismiss her as a credible source on the basis of the fact that she doesn't fit into our stereotype of 'car mechanic'. This dismissal may be more or less conscious or controllable (Levy 2015).

One question is what kind of a mistake people are making when they rely on these kinds of cues, consciously or unconsciously. It is not obvious how individuals can resist the pervasive influence of cultural stereotypes, especially if stereotypes are generally inculcated by statistical learning processes which often result in true beliefs (Begby 2013; Gendler 2011; Egan 2011). In some cases, skewed media presentations of a group may inculcate stereotypes which are frankly inaccurate, for example associating Black Americans with guns, when they are only half as likely as White Americans to be gun owners (Pew Research Survey, 2013). But even in cases where the stereotype is accurate—assume, for the sake of argument, that women know less about cars than men do—discounting the testimony of an individual female car mechanic on the basis of such a stereotype seems morally problematic. Epistemic evaluations are deeply linked to relationships of social trust, and have broad implications for how we treat others and

learn from them. Closer psychological and philosophical examination of the nature of these evaluations can help us to appreciate their power, understand their shortcomings, and search for ways to correct them.

Bibliography

- Aichhorn, M. et al., 2009. Temporo-parietal junction activity in theory-of-mind tasks: falseness, beliefs, or attention. *Journal of Cognitive Neuroscience*, 21(6), pp.1179–1192.
- Apperly, I.A. et al., 2004. Frontal and temporo-parietal lobe contributions to theory of mind: neuropsychological evidence from a false-belief task with reduced language and executive demands. *Journal of Cognitive Neuroscience*, 16(10), pp.1773–1784.
- Apperly, I.A. et al., 2006. Intact first- and second-order false belief reasoning in a patient with severely impaired grammar. *Social Neuroscience*, 1(3–4), pp.334–348.
- Apperly, I.A., 2011. *Mindreaders: The Cognitive Basis of “Theory of Mind,”* Hove and New York: Psychology Press.
- Armstrong, D., 1969. Does Knowledge Entail Belief? *Proceedings of the Aristotelian Society*, 70, pp.21–36.
- Bach, K., 2005. The Emperor’s New “Knows.” In G. Preyer & G. Peter, eds. *Contextualism in Philosophy: Knowledge, Meaning, and Truth*. New York: Oxford University Press, pp. 51–89.
- Baillargeon, R., Scott, R.M. & He, Z., 2010. False-belief understanding in infants. *Trends in cognitive sciences*, 14(3), pp.110–118.
- Begby, E., 2013. The Epistemology of Prejudice. *Thought: A Journal of Philosophy*, 2(2), pp.90–99.
- Behrens, T.E.J. et al., 2008. Associative Learning of Social Value. *Nature*, 456(7219), pp.245–250.
- Birch, S. & Bloom, P., 2004. Understanding children’s and adults’ limitations in mental state reasoning. *Trends in Cognitive Sciences*, 8(6), pp.255–260.
- Bond, C.F. & DePaulo, B.M., 2006. Accuracy of Deception Judgments. *Personality and Social Psychology Review*, 10(3), pp.214–234.
- Buckwalter, W., Rose, D. & Turri, J., 2015. Belief through Thick and Thin. *Noûs*, 49(4), pp.748–775.
- Buttelmann, D., Carpenter, M. & Tomasello, M., 2009. Eighteen-month-old infants show false belief understanding in an active helping paradigm. *Cognition*, 112(2), pp.337–342.
- Butterfill, S.A. & Apperly, I.A., 2013. How to construct a minimal theory of mind. *Mind & Language*, 28(5), pp.606–637.
- Carrington, S.J. & Bailey, A.J., 2009. Are there theory of mind regions in the brain? A review of the neuroimaging literature. *Human Brain Mapping*, 30(8), pp.2313–2335.
- Carruthers, P., 2016. Two Systems for Mindreading? *Review of Philosophy and Psychology*, 7(1), pp.141–162.
- Catmur, C. et al., 2016. Avatars and Arrows in the Brain. *NeuroImage*, 132, pp.8–10.
- Chow, V., Poulin-Dubois, D. & Lewis, J., 2008. To see or not to see: Infants prefer to follow the gaze of a reliable looker. *Developmental Science*, 11(5), pp.761–770.
- Clément, F., 2010. To Trust or not to Trust? Children’s Social Epistemology. *Review of Philosophical Psychology*, 1, pp.531–549.
- Clément, F., Koenig, M.A. & Harris, P., 2004. The Ontogenesis of Trust. *Mind & Language*, 19(4), pp.360–379.
- Clements, W.A. & Perner, J., 1994. Implicit Understanding of Belief. *Cognitive Development*, 9(4), pp.377–395.
- Cogsdill, E.J. et al., 2014. Inferring Character From Faces: A Developmental Study. *Psychological Science*, 25(5), pp.1132–1139.
- Cohen, S., 1999. Contextualism, Skepticism, and the Structure of Reasons. *Philosophical Perspectives*, 13(s13), pp.57–89.
- Couch, J. V. & Sigler, J.N., 2002. Gender of an Expert Witness and the Jury Verdict. *The Psychological Record*, 52, pp.281–287.
- Davies, M. & Gardner, D., 2010. *Frequency Dictionary of American English*, New York: Routledge.
- DeRose, K., 2009. *The Case for Contextualism: Knowledge, Skepticism, and Context, Volume 1*, New York: Oxford University Press.
- Egan, A., 2011. Comments on Gendler’s “The Epistemic Costs of Implicit Bias.” *Philosophical Studies*, 156(1), pp.65–79.

- Egre, P., 2008. Question-Embedding and Factivity. *Grazer Philosophische Studien*, 77(1), pp.85–125.
- Ekman, P. & O’Sullivan, M., 1991. Who Can Catch a Liar? *American Psychologist*, 46(9), pp.913–920.
- Fantl, J., 2015. What Is It to Be Happy That P? *Ergo, an Open Access Journal of Philosophy*, 2.
- Fiske, S.T., Cuddy, A.J.C. & Glick, P., 2006. Universal Dimensions of Social Cognition: Warmth and Competence. *Trends in Cognitive Sciences*, 11(2), pp.77–83.
- Fricker, E., 2009. Is Knowing a State of Mind? The Case Against. In P. Greenough & D. Pritchard, eds. *Williamson on Knowledge*. New York: Oxford University Press, pp. 31–60.
- Fricker, M., 2007. *Epistemic Injustice: Power and the Ethics of Knowing*, Oxford: Oxford University Press.
- Frith, U. & Frith, C.D., 2003. Development and Neurophysiology of Mentalizing. *Philosophical Transactions of the Royal Society of London. Series B, Biological sciences*, 358(1431), pp.459–73.
- Gendler, T.S., 2011. On the Epistemic Costs of Implicit Bias. *Philosophical Studies*, 156(1), pp.33–63.
- Gerken, M., 2013. Epistemic Focal Bias. *Australasian Journal of Philosophy*, 91(1), pp.41–61.
- Gettier, E.L., 1963. Is Justified True Belief Knowledge? *Analysis*, 23(6), pp.121–123.
- Gilbert, D.T. & Malone, P.S., 1995. The Correspondence Bias. *Psychological Bulletin*, 117(1), pp.21–38.
- Goldman, A., 1994. Naturalistic Epistemology and Reliabilism. *Midwest Studies in Philosophy*, 19(1), pp.301–320.
- Harris, P.L., Yang, B. & Cui, Y., 2016. “I don’t know”: Children’s early talk about knowledge. *Mind & Language*.
- Heintz, C. et al., 2010. Epistemic Vigilance. *Mind & Language*, 25(4), pp.359–393.
- Helming, K.A., Strickland, B. & Jacob, P., 2014. Making Sense of Early False-Belief Understanding. *Trends in Cognitive Sciences*, 18(4), pp.167–170.
- Heyes, C., 2014. False Belief in Infancy: A Fresh Look. *Developmental Science*, 17(5), pp.647–659.
- Kaminski, J., Call, J. & Tomasello, M., 2008. Chimpanzees Know What Others Know, but Not What They Believe. *Cognition*, 109(2), pp.224–234.
- Karttunen, L., 1977. Syntax and Semantics of Questions. *Linguistics and Philosophy*, 1(1), pp.3–44.
- Keysar, B., Lin, S. & Barr, D.J., 2003. Limits on Theory of Mind Use in Adults. *Cognition*, 89(1), pp.25–41.
- Kim, M. & Yuan, Y., 2015. No Cross-Cultural Differences in Gettier Car Case Intuition: A Replication Study of Weinberg et al. 2001. *Episteme*, 12(3), pp.355–361.
- King, J.C., Soames, S. & Speaks, J., 2014. *New Thinking About Propositions*, Oxford: Oxford University Press.
- Kiparsky, P. & Kiparsky, C., 1970. Fact. In M. Bierwisch & K. E. Heidolph, eds. *Progress in Linguistics*. The Hague: Mouton, pp. 143–173.
- Koenig, M.A. & Harris, P.L., 2005. Preschoolers mistrust ignorant and inaccurate speakers. *Child Development*, 76(6), pp.1261–1277.
- Koenig, M.A. & Woodward, A.L., 2010. Sensitivity of 24-Month-Olds to the Prior Inaccuracy of the Source: Possible Mechanisms. *Developmental Psychology*, 46(4), pp.815–826.
- Kuhlmeier, V., Wynn, K. & Bloom, P., 2003. Attribution of Dispositional States By 12-Month-Olds. *Psychological Science*, 14(5), pp.402–408.
- Lahiri, U., 2002. *Questions and Answers in Embedded Contexts*, Oxford: Oxford University Press.
- Lee, M.D., Newell, B.R. & Vandekerckhove, J., 2014. Modeling the adaptation of search termination in human decision making. *Decision*, 1(4), pp.223–251. Available at: <http://doi.apa.org/getdoi.cfm?doi=10.1037/dec0000019>.
- Lehrer, K., 1968. Belief and Knowledge. *The Philosophical Review*, 77(4), p.491.
- Levy, N., 2015. Neither Fish nor Fowl: Implicit Attitudes as Patchy Endorsements. *Nous*, 49(4), pp.800–823.
- Lewis, D., 1996. Elusive Knowledge. *Australasian Journal of Philosophy*, 74(4), pp.549–567.
- Luo, Y. & Baillargeon, R., 2007. Do 12.5-month-old infants consider what objects others can see when interpreting their actions? *Cognition*, 105(3), pp.489–512.
- Machery, E. et al., 2015. Gettier Across Cultures. *Nous*.
- Martcorena, D.C.W. et al., 2011. Monkeys Represent Others’ Knowledge but not Their Beliefs. *Developmental Science*, 14(6), pp.1406–1416.
- McKimmie, B.M. et al., 2004. Jurors’ Responses to Expert Witness Testimony: The Effects of Gender Stereotypes. *Group Processes & Intergroup Relations*, 7(2), pp.131–143.
- Milligan, K., Astington, J.W. & Dack, L.A., 2007. Language and Theory of Mind: Meta-Analysis of the Relation Between Language Ability and False-Belief Understanding. *Child Development*, 78(2), pp.622–646.
- Myers-Schulz, B. & Schwitzgebel, E., 2013. Knowing that P without believing that P. *Nous*, 47(2), pp.371–384.
- Nagel, J., 2010a. Epistemic Anxiety and Adaptive Invariantism. *Philosophical Perspectives*, 24, pp.407–435.
- Nagel, J., 2010b. Knowledge ascriptions and the psychological consequences of thinking about error. *Philosophical Quarterly*, 60(239), pp.286–306.
- Nagel, J., Juan, V.S. & Mar, R.A., 2013. Lay Denial of Knowledge for Justified True Beliefs. *Cognition*, 129(3),

- pp.652–661.
- Neal, T.M.S., 2014. Women as Expert Witnesses: A Review of the Literature. *Behavioral Sciences & The Law*, 32, pp.164–179.
- O’Sullivan, M., 2003. The Fundamental Attribution Error in Detecting Deception: The Boy-Who-Cried-Wolf Effect. *Personality and Social Psychology Bulletin*, 29(10), pp.1316–1327.
- Onishi, K.H. & Baillargeon, R., 2005. Do 15-month-old Infants Understand False Beliefs? *Science*, 308(5719), p.255.
- Pascual, B. et al., 2008. Acquisition of mental state language in Spanish children: A longitudinal study of the relationship between the production of mental verbs and linguistic development. *Developmental Science*, 11(4), pp.454–466.
- Pasquini, E.S. et al., 2007. Preschoolers monitor the relative accuracy of informants. *Developmental psychology*, 43(5), pp.1216–1226.
- Perner, J., 1991. *Understanding the Representational Mind*, Cambridge, MA: MIT Press.
- Perner, J. & Roessler, J., 2012. From Infants’ to Children’s Appreciation of Belief. *Trends in Cognitive Sciences*, 16(10), pp.519–525.
- Pew Research Survey, Gun Ownership Trends and Demographics.
- Plato, 1990. *The Theaetetus of Plato*, Indianapolis: Hackett.
- Pratt, C. & Bryant, P., 1990. Young Children Understand That Looking Leads to Knowing (So Long as They Are Looking into a Single Barrel). *Child Development*, 61(4), pp.973–982.
- Radford, C., 1966. Knowledge: By Examples. *Analysis*.
- Rakoczy, H., 2015. In defense of a developmental dogma: children acquire propositional attitude folk psychology around age 4. *Synthese*, pp.1–19.
- Reber, R. & Unkelbach, C., 2010. The Epistemic Status of Processing Fluency as Source for Judgments of Truth. *Review of Philosophy and Psychology*, pp.1–19.
- Rose, D. & Schaffer, J., 2013. Knowledge entails dispositional belief. *Philosophical Studies*, 166(SUPPL1), pp.19–50.
- Ross, L., 1977. The Intuitive Psychologist and His Shortcomings: Distortions in the Attribution Process. *Advances in experimental social psychology*, 10, pp.173–220.
- Ruffman, T., 2014. To Belief or not Belief: Children’s Theory of Mind. *Developmental Review*, 34(3), pp.265–293.
- Sabbagh, M.A. & Baldwin, D.A., 2001. Learning Words from Knowledgeable versus Ignorant Speakers: Links Between Preschoolers’ Theory of Mind and Semantic Development. *Child Development*, 72(4), pp.1054–1070.
- Samson, D. et al., 2010. Seeing it Their Way. *Journal of Experimental Psychology-Human Perception and Performance*, 36(5), pp.1255–1266.
- Santesteban, I. et al., 2014. Avatars and Arrows: Implicit mentalizing or domain-general processing? *Journal of Experimental Psychology: Human Perception and Performance*, 40(3), p.929.
- Saxe, R. & Powell, L.J., 2006. It’s the Thought That Counts: Specific Brain Regions for One Component of Theory of Mind. *Psychological Science*, 17(8), pp.692–699.
- Schurz, M. et al., 2015. Clarifying the role of theory of mind areas during visual perspective taking: Issues of spontaneity and domain-specificity. *NeuroImage*, 117, pp.386–396.
- Schurz, M. et al., 2014. Fractionating theory of mind: A meta-analysis of functional brain imaging studies. *Neuroscience & Biobehavioral Reviews*, 42, pp.9–34.
- Seyedsayamdost, H., 2015. On Normativity and Epistemic Intuitions: Failure to Detect Differences between Ethnic Groups. *Episteme*, 12(1), pp.95–116.
- Shatz, M., Wellman, H.M. & Silber, S., 1983. The acquisition of mental verbs: A systematic investigation of the first reference to mental state. *Cognition*, 14(3), pp.301–321.
- Shope, R.K., 1983. *The Analysis of Knowing: A Decade of Research*, Princeton: Princeton University Press.
- Sobel, D.M. & Kushnir, T., 2013. Knowledge Matters: How Children Evaluate the Reliability of Testimony as a Process of Rational Inference. *Psychological Review*, 120(4), pp.1–19.
- Sporer, S.L. & Schwandt, B., 2007. Moderators of Nonverbal Indicators of Deception: A Meta-Analytic Synthesis. *Psychology, Public Policy, and Law*, 13(1), pp.1–34.
- Sporer, S.L. & Schwandt, B., 2006. Paraverbal Indicators of Deception: A Meta-Analytic Synthesis. *Applied Cognitive Psychology*, 20, pp.421–446.
- Stanley, J., 2005. *Knowledge and Practical Interests*, Oxford: Oxford University Press.
- Stephens, E., Suarez, S. & Koenig, M., 2015. *Early Testimonial Learning: Monitoring Speech Acts and Speakers* 1st ed., Elsevier Inc.
- Steup, M. & Ichikawa, J., 2008. The analysis of knowledge. *Stanford Encyclopedia of Philosophy*, <http://plato.stanford.edu/archives/win2008/entries/knowledge/>
- Stine, G.C., 1976. Skepticism, Relevant Alternatives, and Deductive Closure. *Philosophical Studies*, 29(4), pp.249–261.
- Tardif, T. & Wellman, H.M., 2000. Acquisition of mental state language in Mandarin-and Cantonese-speaking

- children. *Developmental Psychology*, 36(1), p.25.
- Tummeltshammer, K.S. et al., 2014. Infants Track the Reliability of Potential Informants. *Psychological Science*, 25(9), pp.1730–1738.
- Turri, J., 2016. A New Paradigm for Epistemology: From Reliabilism to Abilism. *Ergo*.
- Vanderbilt, K.E., Heyman, G.D. & Liu, D., 2014. In the absence of conflicting testimony young children trust inaccurate informants. *Developmental Science*, 17(3), pp.443–451.
- Varley, R., Siegal, M. & Want, S.C., 2001. Severe Impairment in Grammar Does Not Preclude Theory of Mind. *Neurocase*, 7(6), pp.489–493.
- Weatherston, B., 2005. Can We Do Without Pragmatic Encroachment? *Philosophical Perspectives*, 19(1), pp.417–443.
- Weinberg, J., Nichols, S. & Stich, S., 2001. Normativity and Epistemic Intuitions. *Philosophical Topics*, 29(1), pp.429–460.
- Wellman, H.M., Cross, D. & Watson, J., 2001. Meta Analysis of Theory of Mind Development: The Truth About False Belief. *Child Development*, 72(3), pp.655–684.
- Williamson, T., 2005. Contextualism, Subject-Sensitive Invariantism and Knowledge of Knowledge. *Philosophical Quarterly*, 55(219), pp.213–235.
- Williamson, T., 2000. *Knowledge and its Limits*, New York: Oxford University Press.
- Willis, J. & Todorov, A., 2006. First Impressions: Making Up Your Mind After a 100-ms Exposure to a Face. *Psychological Science*, 17(7), pp.592–598.
- Wimmer, H. & Perner, J., 1983. Beliefs about Beliefs: Representation and Constraining Function of Wrong Beliefs in Young Children's Understanding of Deception. *Cognition*, 13(1), pp.103–128.
- Woolley, J.D. & Ghossainy, M.E., 2013. Revisiting the Fantasy-Reality Distinction: Children as Naïve Skeptics. *Child Development*, 84(5), pp.1496–1510.
- Young, L. et al., 2010. Disruption of the right temporoparietal junction with transcranial magnetic stimulation reduces the role of beliefs in moral judgments. *Proceedings of the National Academy of Sciences*, 107(15), pp.6753–6758.
- Zebrowitz, L.A., Voinescu, L. & Collins, M.A., 1996. “Wide-Eyed” and “Crooked-Faced”: Determinants of Perceived and Real Honesty Across the Lifespan. *Personality and Social Psychology Bulletin*, 22(12), pp.1258–1269.