

- Ratcliff R (1990) Connectionist models of recognition memory: constraints imposed by learning and forgetting functions. *Psychological Review* 97: 285–308.
- Ratcliff R, Van Zandt T and McKoon G (1995) Process dissociation, single-process theories, and recognition memory. *Journal of Experimental Psychology: General* 124: 352–374.
- Ratcliff R and McKoon G (2000) Memory models. In: Tulving E and Craik FIM (eds) *The Oxford Handbook of Memory*, pp. 571–581. New York, NY: Oxford University Press.
- Rolls ET and Treves A (1998) *Neural Networks and Brain Function*. New York, NY: Oxford University Press.

# Epistemology

Intermediate article

Joseph Cruz, Williams College, Williamstown, Massachusetts, USA

## CONTENTS

*What is epistemology?*

*Epistemology and cognitive science*

*Epistemology is the philosophical study of what is required in order to have rational beliefs and knowledge. Both traditional a priori methods of philosophy and a posteriori methods of cognitive science have been brought to bear on this question.*

## WHAT IS EPISTEMOLOGY?

Epistemology answers to a daunting variety of senses in the humanities and the social sciences. Even when we restrict our attention to epistemology as it is understood in contemporary Anglo-American philosophy, the only uncontroversial claim we can make is that epistemology is an attempt to make sense of the possibility, nature, and limits of human intellectual achievement. Typically, the epistemologist does this by trying to illuminate the difference between knowledge and opinion, or the difference between good reasoning and poor reasoning. This project is distinct from merely giving a descriptive account of what people claim to know or to believe reasonably. Instead, epistemologists try to understand what it is really to know or really to believe reasonably, even if people routinely fail to know or are frequently irrational. Moving beyond the descriptive details of knowledge or belief formation to what people ought to believe is a normative philosophical enterprise.

Construed one way, epistemology aims to understand general and ubiquitous elements of human inquiry, such as perceptual knowledge or inductive inference. This project has sometimes

been fueled by skeptical doubts about the veracity of our senses or the trustworthiness of our reasoning. Not all philosophers are persuaded, however, that thoroughgoing skepticism allows for, or requires, a response. As a result we often find in contemporary epistemology the attempt to account for epistemic achievements in a way that does not necessarily offer a reply to the skeptic.

Construed another way, epistemology aims to investigate specific domains of knowledge or rational belief. Some aspects of the philosophy of science may thus be understood as constituting a subfield of epistemology. This kind of research may be narrowed to particular sciences, such as the philosophy of psychology or the philosophy of cognitive science (or physics, or biology, to name two more prominent areas). Efforts to understand the nature of explanation in cognitive science ultimately fall under the umbrella of epistemology, though such efforts lie at the intersection of philosophy of science and philosophy of mind and have thereby taken on a robust theoretical autonomy.

Another distinction may be drawn between epistemology oriented towards individuals and epistemology oriented towards social institutions or practices. Some philosophers claim that there are social practices that positively or negatively influence the formation of knowledge or rational belief. A strong version of this view is where knowledge or rationality is exhausted by socially-mediated factors. This position has currency in some fields of the humanities, but the epistemologists who interact most frequently with the cognitive sciences

tend to reject the most radical forms of social epistemology.

We shall concentrate on the general construal of individual epistemology and its relation to cognitive science.

## Knowledge

What is knowledge? One sort of knowledge is the kind expressed by 'that' clauses. For instance, we may say 'Carlos knows that Oaxaca is in Mexico'. This kind of knowledge is called propositional knowledge, as it is a proposition – in this case, 'Oaxaca is in Mexico' – that is known. Another kind of knowledge is that expressed by 'how to' clauses. Thus, we may say 'Carlos knows how to ride a bicycle'. This second kind of knowledge is called procedural or nonpropositional knowledge. The bulk of attention in contemporary epistemology has been on propositional knowledge, but this emphasis should not be viewed as making a claim about which kind is more important in human cognition.

Beliefs are thought to be the primary psychological entities that are candidates for propositional knowledge. Of course, not every belief is an instance of knowledge. Carlos may believe that Oaxaca is in Mexico, but his merely believing it does not seem to be enough to make it an instance of knowledge. Carlos cannot know that Oaxaca is in Mexico if it is false that Oaxaca is in Mexico. So one additional thing required for a belief to be knowledge is that the belief be true.

At least since Plato, epistemologists have thought that true belief, while necessary, is still insufficient for knowledge. One way of appreciating this is to note that beliefs that are accidentally true are not knowledge. Suppose that Carlos is guessing that Oaxaca is in Mexico. In that case, his belief does not seem to count as an instance of knowledge. One prominent and historically important way of pursuing the distinction between accidentally true and nonaccidentally true belief is to rely on the goodness of the reasons for the belief. So, if Carlos' belief that Oaxaca is in Mexico is both true and based on his having visited that state, we are likely to have a case of knowledge (assuming, for this example, that his visiting Oaxaca is a source of good reasons). On this conception of knowledge three central projects of epistemology emerge. First, the epistemologist must determine what constitute reasons. Second, she or he must offer a general account of what makes some reasons good. And third, the epistemologist must illuminate the nature of the relationship between reasons and beliefs.

Unfortunately, the 'goodness of reasons' approach to knowledge seems susceptible to a curious kind of problem that shows that true beliefs for which we have good reasons still might not count as knowledge. I may believe and have very good reasons to believe, for instance, that 'Mary owns a Honda'. I could then, for the purposes of teaching a logic course, create a variety of compound sentences involving my belief that Mary owns a Honda. Suppose I propose that 'Mary owns a Honda OR Sally is in Paris' even though I do not have any idea where Sally is. In the normal case, if I know one half of this disjunctive compound sentence, then I know the entire sentence since arbitrary disjuncts cannot change the truth-value of the entire sentence no matter how improbable they are. Now imagine that, contrary to my excellent evidence, Mary does not own a Honda, and, completely coincidentally, Sally *is* in Paris. It then seems that my belief in the whole expression 'Mary owns a Honda or Sally is in Paris' is true (Sally's being in Paris makes it so) and is something that I have good evidence for (since I have good evidence that Mary owns a Honda even though she does not). The problem is that we would not be inclined to attribute knowledge to me, though the three conditions for knowledge – true belief with good reasons – have been satisfied.

This is called the Gettier problem, after the author of the short article that sparked contemporary interest in it (Gettier, 1963). The Gettier problem has inspired many putative solutions, counterformulations, and modified putative solutions. There is no agreement on how the Gettier problem is to be solved within the goodness of reasons approach, but many epistemologists think that what the problem reveals is that there must be some fourth condition to ensure that truth and good reasons will be tied together in knowledge. Alternatively, the puzzle may show that the concept of knowledge is ill-defined. The Gettier problem has by no means crippled epistemology. Even if knowledge is not a concept that epistemologists can characterize in any simple way, issues in epistemology having to do with the nature of reasons and belief remain.

We shall turn to these issues shortly, but first it should be noted that there are accounts of knowledge that do not rely on the goodness of reasons approach. Some of these are pursued because they seem to offer a way to avoid the Gettier problem, while others seem to hold greater promise in defeating skepticism. We can divide these accounts of knowledge that do not rely on the goodness of reasons into two categories. One is where beliefs

are calibrated to the truth in a way that can be characterized by epistemologists even though the relation between belief and truth may not involve reasons (Nozick, 1981). The knower does not have knowledge by virtue of having reasons for a belief; rather, a belief is an instance of knowledge if certain metaphysical or logical facts about the relation between belief and truth are satisfied.

The second category is *contextualism*, which maintains that the truth of attributions of propositional knowledge vary from context to context (DeRose, 1995). For instance, we may properly claim that Carlos knows that Oaxaca is in Mexico in mundane conversational contexts while the same claim would be improper in the rarified context of a university seminar on skepticism. Developing contextualist theories of knowledge requires an elaboration of what makes some contexts more demanding than others as well as a systematic treatment of changes in context.

### Justification, Rationality, and Warrant

So far we have been identifying epistemology with the elucidation of knowledge. Being an instance of knowledge, however, is not the only epistemically positive characteristic a belief might have. We may be interested in having good reasons for our beliefs without insisting that such beliefs be instances of knowledge. The possibility of independently exploring epistemically positive beliefs that fall short of knowledge was implicit above when the three central projects regarding reasons (their nature, what makes some good, and what their relationships are) were proposed. In spite of their potential independence from knowledge, though, all three of these projects are typically conceived against the background of treating truth as the fundamental aim of epistemic reasoning. Good reasons are therefore often understood in the first instance to be reasons in favor of taking a belief to be true. Maintaining a close link between the goodness of reasons approach and truth is a deep commitment in epistemology, and explains the confidence in the connection between rationality and knowledge. Still, some epistemologists have offered accounts that sever the link by treating good reasoning as wholly unconnected to the truth.

The most prominent properties of epistemically positive beliefs discussed by philosophers are justification and rationality. These labels carry with them some connotative differences. Except where the distinction is crucial, in this essay we will read 'justification' and 'rationality' as synonyms. *Warrant* also has some currency in describing

what a belief must have, in addition to truth, in order to yield knowledge.

How should we undertake evaluations of beliefs in terms of justification or rationality or warrant? The goodness of reasons strategy for investigating knowledge again affords a persuasive framework. We may study the reasons for a belief in order to make some judgment about whether that belief is epistemically positive. Furthermore, if rationality or warrant are graded notions, the strategy may allow us to advance a scale that appeals to the comparative goodness of the reasons for a belief. Thus, we may account for judgments of 'more rational' or 'more warranted' belief.

There seems to be an immediate problem, though, with employing the strategy of reasons for assessing epistemically positive belief along the dimensions of justification or rationality. A belief will inherit the epistemic status of the reasons for it, so we must in turn determine whether the reasons for a belief are themselves rational or justified. This threatens a regress of reasons that must be resolved. Responding to this regress has been instrumental in crystallizing issues in the recent history of epistemology, even if the regress problem no longer occupies the most crucial role in contemporary epistemological research. Deflecting the regress of reasons argument does not by itself answer the question of what makes a reason justified or rational. An answer to the structural question seems like a necessary first step, but it cannot be a complete theory of epistemically positive beliefs.

### ***The structure of the belief corpus: foundationalism and coherentism***

Two responses to the regress problem that differ on the structure of the belief corpus are foundationalism and coherentism. The foundationalist claims that there is a set of basic beliefs that do not require reasons to explain their epistemically positive nature because of some special characteristic(s) that they have. The epistemic credentials of beliefs that are not foundational are due to a traceable lineage through reasons, from basic beliefs via a basing relation that must be illuminated by the foundationalist. Thus, the justification relationship in the foundationalist picture is asymmetrical.

Foundationalists have attempted to develop axioms of goodness for foundational beliefs. Candidates for axioms include the claim that beliefs about how things appear are intrinsically epistemically positive. In a like fashion, foundationalists will need to provide enough axioms for foundational beliefs to account for the credentials of all epistemically positive beliefs.

Modulating the claim that the regress of reasons must end in intrinsically epistemically positive beliefs can complicate the foundationalist picture. It is an open question whether reasons for beliefs must themselves be beliefs. If not, then it is possible that the foundational reasons are other cognitive states such as perceptual states or memory states.

In contrast to foundationalists, coherentists maintain that no beliefs are intrinsically epistemically positive (Bonjour, 1985). By their lights, every belief relies on other beliefs for its epistemic status. One (controversial) argument for coherentism is that the foundationalist's putative axioms of goodness require an argument to show that they are good, and any such argument will rely on further beliefs, indicating that the axioms are not foundational after all. On one reading of coherentism, beliefs are epistemically positive based on a lineage of reasons in a structure that may ultimately loop back onto itself. There may be no need, however, to trace reasons in a way that is circular. One might instead claim that a belief is epistemically positive in the case of its being a member of a coherent belief corpus without pursuing particular reasons in a linear fashion.

Coherentists discharge the task of revealing what makes some reasons good by claiming that cogent arguments can be given for the high epistemic credentials of beliefs, and further arguments can be given in favor of the cogent arguments.

### ***Epistemic internalism***

Despite their differences on the structure of the belief corpus, the foundationalist and coherentist strategies traditionally agree on a different issue: reflective, careful agents are able to make assessments of their own beliefs in order to determine whether they are epistemically positive. This is called the *internalist* conception of epistemic justification or rationality.

In order to defend internalism in detail, the epistemologist needs to specify what is meant by the claim that agents can determine whether there are good reasons for their beliefs. A number of interacting but independent issues arise here. One challenge is deciding whether the epistemic agent must be able to determine that reasons are *good*, versus the less strict demand that the agent merely be able to determine what the reasons *are*. Another challenge is whether the agent must be able to determine *what* her or his good reasons are for a belief, versus the much less demanding constraint that the agent must be able to determine *that* she or he has good reasons for a belief. A third challenge is deciding what sort of effort on the part of the agent

will be consistent with the claim that agents can access their reasons. Most of the combinations of answers to these challenges have been defended in the internalist literature.

Internalism originates in three related concerns. First, one of the projects that sometimes accompanies an assessment of epistemically positive belief is to illuminate how one might improve one's beliefs. If improvement is to be possible, it needs to be possible to determine which belief among many candidate beliefs is most epistemically positive, and it has seemed that the epistemic agent personally needs to be able to make the judgment. Second, justification (though not rationality) has often been viewed as at least partly a matter of fulfilling a distinctly epistemic duty. Fulfilling a duty seems to require that one be able to do the things that duty requires. In order to secure the means to an intellectual duty, an epistemic agent will need to be able to reflect on her or his condition and on the resources available. Third, recall that the ability to answer the skeptic is sometimes thought to be a crucial component of epistemology. The only answers that the epistemic agent can give, though, are ones that are available on reflection.

### ***Epistemic externalism***

Thinking of the epistemic agent as able to determine when beliefs are epistemically positive is not demanded by the overarching goal of epistemology. Though the philosopher may state the conditions that must be met for a belief to be epistemically positive, the epistemic agent may not be able to make assessments of particular beliefs. Thus one might reject internalism and not expect the right philosophical account of epistemically positive belief to enable meliorative, duty-oriented, or skeptic-answering evaluations to take place. This is the *externalist* view in epistemology.

For example, *process reliabilism* – the best-known externalist view – to a first approximation claims that what confers positive epistemic status on a belief is that it be produced by a psychological process that reliably produces true belief (Goldman, 1979). Though the reliability of a psychological process is often opaque to the person employing that process, process reliabilists think that the reliability of belief-forming processes can be uncovered by cognitive science or by other kinds of empirical inquiry.

The motivations for defending an externalist position are diverse. One simple motivation is the impulse to tie justification and rationality directly to truth via an appeal to truth-sensitive properties such as reliability. In this connection, the

externalist's commitment to truth as the central goal of justified belief looms large. Externalist theories can be neutral on the question of the structure of the belief corpus, and instead attempt to tackle more directly the issue of epistemically positive belief.

Another motivation that is prominent in discussions of externalism is the role of causal factors in belief formation. It has seemed to many epistemologists that, in order to be justified or rational, a belief has to have both an evidential and a causal relation to the reasons for it. Incorporating this causal element seems to require some of the specialist's insight into the causes of our beliefs, and that is in tension with the internalist's impulse to insist that it is the epistemic agent personally who is in a position to determine the epistemic status of her or his beliefs.

### The Methods of Epistemology

In traditional epistemology, the standard of correctness for epistemological questions appeals to intuitions about epistemic methods or particular cases of belief. For example, that beliefs formed under favorable perceptual conditions in a healthy observer are justified is a principle that we may intuitively certify. So, a non-skeptical account of justification should in this view be designed to accommodate intuitions about beliefs formed under favorable perceptual conditions. Alternatively, a theorist may take particular instances of an intuitively epistemically positive belief – the belief that *this* shiny Macintosh apple is red, for instance – and attempt to build a theory that respects this intuition. A successful theory of justification will yield the result that such beliefs are justified unless there is some other overriding consideration that would result in the retraction of the judgment that the belief is justified. Such beliefs are called 'prima facie justified' in order to highlight the fact that new information about the situation might change intuitions about the justifiedness of the belief in question. Many epistemologists have exploited both kinds of intuition by seeking to balance methods and cases.

There are reasons to be dissatisfied with the methodology of epistemology, not the least of which is that the use of intuitions alone can seem to be a precarious or even spurious basis for theorizing. Intuitions might be thought to be too subjective or theory-laden. There was a period in Anglo-American philosophy when philosophy was viewed as wholly conceptual (*a priori*) analysis. The strong commitment to that position has long since

passed. Intuitions ground philosophical inquiry in our pre-theoretic understanding of epistemic concepts, and allow a bridge to the long history of philosophical theorizing on the same topics. But intuitions may need to be adjusted in light of other considerations. There is much debate as to what the other sources of constraint may be. It is within a particular answer in this debate that there is the most cross-fertilization between epistemology and cognitive science.

## EPISTEMOLOGY AND COGNITIVE SCIENCE

Some epistemologists maintain that a sensible division of labor in understanding justification and rationality is that philosophers investigate the normative elements of belief formation, while cognitive scientists study how we actually form beliefs. To the extent, however, that philosophical intuitions about epistemic concepts need to be constrained by the empirical (*a posteriori*) details of belief formation, cognitive science will play a significant role in epistemology. Theories of rational belief that incorporate a significant empirical component are categorized under the label of 'naturalized epistemology'.

### Naturalized Epistemology

Cognitive science gives epistemologists detailed and empirically robust accounts of the origins of belief. This is relevant to epistemology if we are persuaded that a belief is rational only if it bears the right causal relation to the reasons for it. Empirical details will also be important in part because philosophical intuitions are likely to be misleading with respect to the nuances of belief formation. For example, intuitively it seems that inductive reasoning on the basis of small samples would not be a method for forming epistemically positive beliefs. It has been argued, however, that induction on the basis of small samples is epistemically defensible in reasoning involving natural kinds (Kornblith, 1993). It is only through the insight of cognitive scientific accounts of how reasoning interacts with kind concepts that this unintuitive method will appear epistemologically sound. Other insights from cognitive science that have been relevant in epistemology include vision research, research on deductive reasoning, and research on memory.

Process reliabilism advocates a proprietary role for cognitive science in epistemology. Once causal factors are deemed relevant in understanding

justification, and once it has been determined that justified beliefs are the product of reliable psychological processes, cognitive science enters into epistemology to provide the details of psychological processes in terms of their reliability (Goldman, 1986).

Naturalized epistemology is often identified with externalism. Depending on how issues of access to reasons are resolved, though, it is possible to maintain an internalist theory that carves out a role for empirical research on belief formation. Artificial intelligence models of cognition have been employed in epistemology to investigate how complex patterns of reasoning interact in order to yield intuitively justified beliefs (Pollock and Cruz, 1999; Thagard, 2000). The rules or norms of reasoning at each step of reasoning may be unintuitive owing to their complexity, but the whole reasoning process may be counted as rational because it yields an intuitively justified conclusion. This position can be internalist by maintaining that epistemic agents will be able to determine in their own cases whether their reasoning is good, even though the norms of reasoning generally will be inscrutable.

## The Cognitive Science of Epistemology

Finally, we consider another way to conceive of the relationship between epistemology and cognitive science. Cognitive science may step back from philosophy and investigate epistemology itself. There are two plausible targets of investigation: concepts and methods. First, cognitive science might investigate how people employ the concepts of justification, knowledge, or good evidence in the same way that nonphilosophical concepts are studied (Goldman, 1992). Drawing conclusions from such research is a delicate matter, however, since it has seemed to philosophers that there is no way to uncover the genuine normative character of epistemological concepts by merely studying their psychological status.

Second, cognitive science might investigate the method of employing *a priori* intuitions to address epistemic questions. A specific instance of this suggestion is that we should think of epistemic intuitions as analogous to grammaticality judgments in linguistics (Pollock and Cruz, 1999). Reasoning on this view is the product of an underlying rational competence, but actual cases of reasoning might go wrong due to resource constraints. If the analogy could be successfully developed, it would help

explain the difference between merely describing how people reason (which would always be particular performances of reasoning) and the normative account of how people ought to reason. The explanatory credentials of this account appeal directly to its success in linguistics and related areas of psychology and artificial intelligence research.

## References

- Bonjour L (1985) *The Structure of Empirical Knowledge*. Cambridge, MA: Harvard University Press.
- DeRose K (1995) Solving the skeptical problem. *The Philosophical Review* 104: 1–52.
- Gettier E (1963) Is justified true belief knowledge? *Analysis* 23: 121–123.
- Goldman A (1979) What is justified belief? In: Pappas G (ed.) *Justification and Knowledge*, pp. 1–14. Boston, MA: D. Reidel.
- Goldman A (1986) *Epistemology and Cognition*. Cambridge, MA: Harvard University Press.
- Goldman A (1992) Epistemic folkways and scientific epistemology. In: *Liaisons*, pp. 155–175. Cambridge, MA: MIT Press.
- Kornblith H (1993) *Inductive Inference and Its Natural Ground*. Cambridge, MA: MIT Press.
- Nozick R (1981) *Philosophical Explanations*. Cambridge, MA: Harvard University Press.
- Pollock J and Cruz J (1999) *Contemporary Theories of Knowledge*, 2nd edn. Lanham, MD: Rowman & Littlefield.
- Thagard P (2000) *Coherence in Thought and Action*. Cambridge, MA: MIT Press.

## Further Reading

- Alston W (1989) *Epistemic Justification: Essays in the Theory of Knowledge*. Ithaca, NY: Cornell University Press.
- Bernecker S and Dretske F (eds) (2000) *Knowledge: Readings in Contemporary Epistemology*. New York: Oxford University Press. [This anthology collects many of the most important articles in contemporary epistemology.]
- Dretske F (1981) *Knowledge and the Flow of Information*. Cambridge, MA: MIT Press.
- Foley R (1993) *Working Without a Net: A Study of Egocentric Epistemology*. New York, NY: Oxford University Press.
- Goldman A (1993) *Philosophical Applications of Cognitive Science*, chaps 1 and 2. Boulder, CO: Westview Press.
- Goldman A (1999) *Knowledge in a Social World*. New York, NY: Oxford University Press.
- Harman G (1999) *Reasoning, Meaning, and Mind*. New York, NY: Oxford University Press.
- Lehrer K (2000) *Theory of Knowledge*, 2nd edn. Boulder, CO: Westview Press.
- Plantinga A (1993) *Warrant and Proper Function*. New York, NY: Oxford University Press.