

Do religious beliefs have a place within an ‘epistemically naturalized’ cognitive system?

Graham Wood

This document is the author’s manuscript as accepted for publication. The final publication is available at Springer via <https://doi.org/10.1007/s11841-016-0567-4>

Abstract

This paper offers an account of religious belief using the resources available within naturalized epistemology. The account accepts the basic architecture of an individual’s ‘web of belief’ as articulated by Quine and Ullian, and explores some of the further details of the Quinean account of naturalised epistemology. The paper then updates that account drawing on recent work in psychology and cognitive science. The paper also considers the purpose of an individual’s web of belief. Quine’s web of belief has a very tightly constrained purpose, namely, the prediction of future experience in the light of past experience. But, drawing lessons from the feminist critique of analytic epistemology and the philosophy of science, this paper considers the consequences if the purposes of a web of belief are wider than simply predicting future experience in the light of past experience. Drawing on work from the cognitive science of religion the paper suggests a wider set of purposes could result in religious beliefs existing in a web of belief, even when that web of belief is understood purely from the perspective of naturalised epistemology.

Introduction

The anthropologist Pascal Boyer observes that:

Religious ideas can be called “natural” in (at least) two senses. “Natural” are those aspects of religious ideas which depend on noncultural constraints, like the human genome or the capacities of human brains or the properties of the world humans live in. [And] “natural” can be understood as describing a subjective quality, the fact that certain religious postulates are considered perfectly obvious, self-evident ideas by the people who hold them (1994, p. 3).

Boyer’s observations are representative of the *naturalness of religion* thesis, a thesis being considered by a number of researchers (e.g., McCauley, 2000; Barrett, 2004; McCauley and Cohen 2010) within the emerging field of Cognitive Science of Religion (CSR). Boyer’s references to ‘capacities of human brains’ and ‘the fact that certain religious postulates are considered perfectly obvious, self-evident ideas’ bring to mind the project of Naturalized Epistemology. Willard van Orman Quine famously (or perhaps infamously) advocated the naturalising of epistemology. For Quine, epistemology ‘or something like it, simply falls into place as a chapter of psychology and hence of natural science’ (1969, p. 82). And Hilary Kornblith, a contemporary advocate of Naturalized Epistemology, identifies the goal of a naturalistic theory of knowledge, not to provide an account of our concept of knowledge, but ‘instead to provide an account of a certain natural phenomenon, namely, knowledge itself’ (1999, p. 161). I wish to bring together Boyer’s observation that some religious ideas are considered

‘natural’ by the people who hold them, and Quine’s view that epistemology should be understood as a chapter of psychology, and then address the question: *Do religious beliefs have a place within an ‘epistemically naturalized’ cognitive system?*

This question raises some significant (and contested) issues. Related to Kornblith’s distinction, between seeking an account of the *natural phenomena* rather than the *concept* of knowledge, is the distinction between the descriptive and the normative attitude to ‘knowledge’. Epistemology is traditionally considered to be a normative endeavour. Epistemology is not traditionally understood as the study of what is *assumed to be knowledge (and perhaps labelled as such in natural language)*, but rather it is the study of *knowledge*. The purpose of epistemology, traditionally understood, is to provide a framework in which one can judge between beliefs that have ‘good standing’ in the relevant sense of knowledge and beliefs that have no such standing. For example, Jaegwon Kim takes the *justification* of belief to be essential to epistemology. He is no advocate of naturalised epistemology, arguing that:

If justification drops out of epistemology, knowledge itself drops out of epistemology. For our concept of knowledge is inseparably tied to that of justification. [K]nowledge itself is a normative notion. Quine’s nonnormative, naturalized epistemology has no room for our concept of knowledge. (Kim 1988, p. 389).

So, according to Kim, Naturalized Epistemology is not epistemology. And Kim is not alone. Others argue that the project of naturalising epistemology simply changes the subject from epistemology to something else (e.g., Bonjour, 1994). Similar contestation surrounds the place and nature of ‘truth’ in deliberations about knowledge. Knowledge may bear some relation to ‘true belief’ (although many who consider such matters assume it must be more than that). But then this leads to the question: what is the correct analysis of ‘truth’? Answers vary (see Glanzberg, 2014) but such analysis takes us beyond the scope of this paper. However, even in the face of these concerns and criticisms, some wish to pursue the project of naturalised epistemology. And in so doing acknowledge the significant implications. For example, Philip Kitcher observes that:

naturalism’s emphasis on using our best scientific knowledge within epistemology should extend to the reformulation of the basic epistemological issues. What the cognitive sciences inform us about the character of our representations should be used in posing those issues (1992, p. 81).

So, while acknowledging the concerns about and criticisms of Naturalized Epistemology, I am encouraged by Kitcher’s suggestion. And in the context of our present exploration, I am encouraged specifically to apply the insights from CSR to help understand the character of religious representations. It is in this context that I address the question: *Do religious beliefs have a place within a ‘epistemically naturalized’ cognitive system?* I intend to examine in a number of themes that will be useful in addressing this question directly and helpful in exploring some related and larger themes. But before I proceed to the detail, I will briefly introduce these themes. What follows is a brief description of central themes rather than a formal argument. But it will give an indication of the explorative thread that I will follow in this paper.

Quine’s (and Ullian’s) web of belief and its purpose.

I will outline Willard van Orman Quine and Joseph Ullian’s web of belief account (1978), and in particular I will examine the role of the *purpose* of the web of belief. Quine has a very narrowly defined purpose of the web of belief (or conceptual scheme). This is to

predict the future in the light of past experience (1980, p. 44). But Quine admits that this is only one purpose and that other purposes are possible. But given the purpose Quine defines, he then identifies a number of theoretical virtues that further his purpose and these theoretical virtues are used to update one's web of belief. So in short, once one's purposes are identified, these lead to the articulation of theoretical virtues, and these theoretical virtues are then used to update one's web of belief.

Lessons learnt from the feminist critique of the philosophy of science

I will then look to claims that have been made within the feminist critique of the philosophy of science and apply the lessons learnt there to the issues being addressed here. Helen Longino (1995) and Louise Antony (2002) have highlighted the fact that while it may be assumed by many that science is objective there are problems with that assumption. These critics have claimed that science is not as objective as assumed. Rather, the so-called objectivity of science serves to further purposes that are not explicitly acknowledged in the self-understanding of the scientific process. The critique considers the traditional theoretical virtues of the sciences and claims that these virtues are informed by a set of purposes wider than the pursuit of 'objectivity'. Furthermore, the feminist critics of science claim that biases are inevitably present in the endorsement of the theoretical virtues. So biases are in play within science. But the feminist critique does not then advocate removing these biases (as removal would be impossible, in their view). Rather it advocates the uncovering of biases and then advocates a process of ongoing identification, exploration and engagement with these biases. I suggest, there is an important parallel here to considerations of the nature of a web of belief more generally and the place of religious belief in that web.

Applying the lessons to the purposes of a Quinean web of belief

I suggest the purposes that lead to the nature and content of a person's web of belief are significantly wider than Quine's identified purpose of predicting the future in the light of past experience. And these wider purposes imply a wider set of theoretical virtues (or perhaps one should call this wider set 'intellectual' or 'cognitive' virtues). As a result of using this wider set of virtues (to revise beliefs within the web), we will find beliefs admitted into (or retained within) the web that would not have been admitted or retained based only on the narrower set of Quinean theoretical virtues. I will suggest that the purposes of individual's webs of belief are wider than simply predicting the future in the light of past experience and if that is accepted, and given the existence of other purposes, the place of religious belief in those webs is understandable.

Updating the Quinean web with evolutionary psychology and cognitive science

As mentioned above, Quine himself endorsed the view that epistemology should be seen as a chapter of psychology. So in this section I draw on a number of developments in evolutionary psychology and cognitive science to update our understanding of the structure of a web of belief. In particular, I introduce 'core knowledge systems' (Spelke and Kinzler 2007), and the 'dual systems' approach to reasoning (Kahneman 2011; Evans and Frankish 2009). Also, drawing from the 'religion as evolutionary adaptation' research program within CSR, I discuss the 'supernatural punishment theory' (Johnson and Bering 2009). Finally, I discuss how the claims I make in this paper, can be brought into engagement with theoretical suggestions being made in contemporary cognitive science. In particular, I draw some parallels between the claims I make in this paper and

claims made by Neil van Leeuwen (2014) regarding the need to distinguish between different cognitive attitudes that can (mistakenly) be taken to be the same cognitive attitude (in other words, many cognitive attitudes are all being considered to be the one cognitive attitude, namely 'belief' and distinctions should to be made between them). That concludes the description of the themes that I will go on to discuss in more detail below. But before I move to the detailed discussion I will offer some brief comments in order to locate my claims within a larger philosophical context.

Quine, Sellars and the purposes of belief

This paper assumes that the Quine/Ullian (1978) account of an individual's 'web of belief' is largely correct. Where this paper does go beyond the Quinean picture is in relation to the purposes of a web of belief. Quine's account of the web is built upon the pursuit of a set of very tightly constrained purposes, namely, 'predicting future experience in the light of past experience' (1980, p.44). This paper will consider a wider set of purposes. But this is not something Quine would find heretical. Indeed, he himself admitted that his purpose of predicting future experience in the light of past experience was only one among other 'interests and purposes' (1980, p. 19).¹ Indeed, the acknowledgement of a wider set of interests and purposes was foreshadowed by another advocate of naturalized epistemology, Wilfred Sellars. Sellars introduced the distinction between the manifest image and the scientific image, and while he did describe them as 'rival' images (1964, p.57), he did not seek to discard the manifest image. As Willem deVries (2015), observes:

Sellars thinks that science's aperspectival description of the empirical world is the measure of reality, but he is also committed to the indispensability of the concepts built into the first-person perspective that makes agency possible. If those concepts were not involved in the regulation of our behavior, we would not be *persons* and could not engage in such activities as moral behavior or scientific research. This is why Sellars calls for a *stereoscopic vision* in which the descriptive resources of the sciences are united with the language of individual and community intentions and the dualism of the manifest and scientific images is transcended.

Following Sellars' lead, and building on the architecture of the web of belief provided by Quine/Ullian, I offer an account of the purposes of a web of belief that explains the incorporation of religious beliefs. If my analysis is accepted, then what are the implications? Put simply, my analysis can be used to understand religious belief. Quine stipulated that the purpose of the web of belief is to predict future experience in the light of past experience. But other people may choose a broader set of purposes for their web of belief.² Or, and perhaps more intriguingly, it may be that evolution also has some 'purposes' in play in the context of this question.³ If it is possible at all to understand the purposes of evolution, then our cognitive systems may have purposes shaped by evolution, and while predicting the future may be one of the purposes, it is, I

¹ Intriguingly, and illustrating Quine's instrumentalism very clearly, he acknowledges his willingness to entertain belief in a Creator if such a belief did the appropriate explanatory work: "If I saw indirect explanatory benefit in positing sensibilia, possibilia, spirits, a Creator, I would joyfully accord them scientific status too, on a par with such avowedly scientific posits as quarks and black holes" (1995, 252).

² Indeed, following Sellars, I suggest that it would be impossible to live as a person in a community if the only purposes of our web of belief was to predict the future in the light of past experience.

³ I am using the term 'purpose' in the context of evolution here as short hand for a longer phrase that describes a process of the generation of cognitive structures and processes as outcomes of evolution.

suggest, unlikely that prediction of the future is the only purpose for which our cognitive faculties have evolved.⁴

Future work (not explored in this paper) could involve using the framework offered by Quine to arrive at 'evolutionarily rational' religious belief (Wood 2013, p. 193), and then by exploring alternate analyses of truth (e.g., pragmatic/deflationary accounts) perhaps arrive at an understanding of what might be called religious knowledge, where that is understood as knowledge in the broader sense of 'intellectual virtue' (Zagzebski 1999, p. 109), or 'cognitive virtue' (Kitcher 1992, p. 66). Finally, prompted by a caveat Quine (1980) offers with respect to his own work (that I will refer to again below), I should note that this paper concerns epistemology and not metaphysics. I am making epistemological claims in this paper, not metaphysical claims.

Quine's web of belief and the conceptual scheme of science

A web of belief (Quine & Ullian, 1978) contains the set of beliefs of an individual, and that web can be imagined as a spider's web, that has many interconnected strands, and is tethered to the world at its edges. The web of belief has the purpose of explaining the past and predicting the future (1978, p. 66). But interestingly, Quine and Ullian choose not to characterise the web as explicitly directed toward truth or knowledge. The edges of the web represent experience (or more precisely 'observational statements'), and it is these observational statements to which the internal content of the web is answerable in the 'tribunal of sensory evidence' (1969, p. 89). In other words, the content of the web must be empirically adequate, in that it must accord with experience (or more accurately accord with the observational statements).

If it is the case that the beliefs of the web do not accord with experience, then the details of the internal connections of the web can be changed to bring the beliefs within the web into accord with experience once more. The centre of the web represents one's core beliefs. Quine uses beliefs about logic (such as the law of the excluded middle) as examples of core beliefs.⁵ Beliefs located at neither the centre nor the edge represent one's 'mid-level beliefs'. Importantly, Quine stipulates that *any* belief in the web is subject to revision, even one's beliefs about logic.⁶ Furthermore, Quine's web of belief is subject to 'confirmational holism' whereby it is the whole web that is to accord with experience. As a consequence of this holism, no individual belief is specifically judged wanting if there is a lack of accord with experience. This is simply because other parts of the web can be adjusted to bring the whole web back into accord with experience. However, if one is unwilling to adjust particular parts of the web, in order to bring the whole web back into accord with experience, then other parts of the web should be changed to achieve accord. Finally, and in large part because of confirmational holism, it is possible that multiple substantially different webs can all be in accord with

⁴ It is not the main focus of this paper, but there is a very interesting research program building momentum examining evolutionary process that may have led to religious belief (see for example, Atran 2002; Schloss & Murray 2009; Wilson 2002). And as an aside, although religious belief may be thought of as making predictions (post mortem predictions, for example), I don't see prediction as a central purpose (if it is a purpose at all) of religious belief.

⁵ But religious believers may also locate certain religious beliefs at the centre of their web.

⁶ And the religious Quinean would similarly have to accept the possibility of the revision of core religious beliefs.

experience. This possibility is referred to as the underdetermination of theory by data (1980), where 'theory' is the content of the web and 'data' is experience. Having introduced the web of belief in general terms, I will now consider some of Quine's more specific claims in the hope of bringing out more of the detail of his position. At this point I will move from talk of a web of belief to talk of the conceptual scheme of science. Now one may be concerned that in this move I have changed the topic of discussion. But Quine himself characterises scientific beliefs as part of the larger web of belief, so I have not changed topic. I have simply focused in on one part of the web of belief, namely scientific beliefs. And I will begin by asking a simple question: *what is the purpose of this conceptual scheme?*

Quine answers this question as follows:

As an empiricist I continue to think of the conceptual scheme of science as a tool, ultimately, for predicting future experience in the light of past experience. (1980, p.44)

So for Quine the conceptual scheme of science is a tool for predicting future experience. This raises a further question: *would the conceptual scheme of science be the best tool for a different job?* And Quine acknowledges that it would not. He describes the question of which conceptual scheme to adopt as an open question, and while he himself addresses it with reference to 'the epistemological point of view', he acknowledges that this point of view is 'one among various, corresponding to one among our various interests and purposes.' (1980, p.19). So Quine acknowledges that if predicting future experience was not the (only) purpose of one's conceptual scheme, then presumably the conceptual scheme of science would not necessarily be the best scheme to adopt. Later in this paper I will return to the question of which conceptual scheme is best suited to broader human interests and purposes and suggest that it may not be the austere scheme Quine advocates for his own interests and purposes.

According to Quine, at the edge of our conceptual scheme (our web) are located our experiences, or more precisely 'observational statements' that relate to those experiences. And 'an observation sentence is one on which all speakers of the language give the same verdict when given the same concurrent stimulation.' (1969, p. 86-87) Importantly, for my analysis, Quine assumes that thought and talk of physical bodies is not *hard-wired* into human experience. He observes that 'we are prone to talk and think of objects', however, it is not the case that 'our objectifying pattern is an invariable trait of human nature' (1969, p.1). Elsewhere he states that observational sentences 'will usually be about bodies' and that 'a corporeal subject matter is likelier than not' (1969, p. 87). And in a famous passage from *Two Dogmas of Empiricism* he writes:

'Physical objects are conceptually imported into the situation as convenient intermediaries -- not by definition in terms of experience, but simply as irreducible posits comparable, epistemologically, to the gods of Homer. ... Both sorts of entities enter our conception only as cultural posits. The myth of physical objects is epistemologically superior to most in that it has proved more efficacious than other myths as a device for working a manageable structure into the flux of experience.' (1980, p.44).

As I alluded to in my introduction, Quine offered a caveat in relation to his 'likening the physicists' posits to the gods of Homer' in his Foreword to the revised second edition of *From a Logical Point of View* (1980, p. viii). There he states that he is concerned with epistemology not metaphysics. It is interesting, in and of itself, that Quine saw the need

to make this statement in 1980. But, for the purposes of this paper, the importance of this qualification relates to Quine's own position on the source of our beliefs about mid-sized physical objects. Quine thinks the source of these beliefs are cultural contingencies. But, as I will suggest below, I think our beliefs about mid-sized physical objects are more *hard-wired* than Quine assumed.

In *Ontological Relativity*, Quine draws our attention to the implicit maxim that guides a linguist's choice to use the term 'rabbit' as referring to a homogenous object is moving against a contrasting background. And with reference to that maxim, he writes:

If he were to become conscious of this maxim he might celebrated it as one of the linguistic universals, or traits of all languages, and he would have no trouble pointing out its psychological plausibility. But he would be wrong; the maxim is his own imposition, toward settling what is objectively indeterminate. It is a very sensible imposition, and I would recommend no other. But I am making a philosophical point. (1969, p.34)

I interpret Quine's position on the status of talk and thought about mid-sized physical objects as 'convenient intermediaries' and as 'cultural posits' as reflecting his empiricist outlook that rejects 'innate ideas', and rejects the assumption that representations of physical objects are universal aspects of human cognition. However, given that Quine's own position is that epistemology is a chapter of psychology it is appropriate to look to recent work in cognitive science to update our understanding of the place of innate ideas in our minds. And, as it happens, evidence has been emerging that suggests innate ideas play a greater role than Quine assumed. But before we consider these issues let us take a brief look at the feminist critique of the philosophy of science. Why, might you ask, is a feminist critique of the philosophy of science relevant to our current discussion? The answer is that if you assume philosophy is continuous with science (and that the conceptual scheme of science is simply part of a wider web of belief), then the feminist critique of the philosophy of science is relevant to epistemology, and relevant to naturalised religious belief.

The feminist critique of the philosophy of science

Feminists have criticised the epistemology of science (Longino 1995; Antony 2002). And I suggest a naturalized religious epistemology can learn from this feminist critique.⁷ In *Quine as Feminist: The Radical Import of Naturalised Epistemology* Antony claims that many feminists think mainstream epistemological paradigms within the analytic tradition are incompatible with feminism (2002, p. 111). However, she goes on to claim that Quine's naturalised epistemology is compatible with feminism, and her goal in her paper is to 'highlight the virtues, from a feminist point of view, of naturalised epistemology' (2002, p. 113). My purpose in this paper is to show that naturalised epistemology may have similar virtues from a religious point of view.⁸

⁷ I should say that I am not an expert in feminist critiques of science in general, nor an expert in the dimensions of the feminist critique that I will be drawing on to make my observations in this paper. But, I hope that I have understood the critique well enough to productively apply it here.

⁸ Interestingly, Plantinga has offered a naturalized account of epistemology and concludes that such an account 'flourishes much better in the garden of supernatural theism than in that of metaphysical or theological naturalism.' (1995, p. 396).

Antony notes that a central feminist criticism of 'mainstream' epistemology concerns the notion of 'objectivity' as an epistemic ideal. Mainstream epistemology, so the criticism goes, has an unrealistic expectation that humans will be able to achieve objectivity. She observes that Quine, Goodman, Hempel, Putman and Boyd have all argued that the received concept of objectivity is untenable as an ideal of epistemic practice (2002, p. 114-115). She explains that the background of this critique of objectivity originated in the disagreement between rationalists and empiricists about the extent of the internal structure of the mind:

The rationalist believed that native structure placed quite specific limitations on the kinds of concepts and hypotheses the mind could form in response to experience, so that human beings were, in effect, natively biased toward certain ways of conceiving the world. Empiricists, on the other hand, held that there were relatively few native constraints on how the mind could organise sensory experience, and that such constraints as did exist were *domain-general* and *content-neutral*. (2002, p. 123).

Both rationalists and empiricists, she continues, believed in:

the possibility of providing a *rational* justification for the processes by which human beings arrive at theories of the world. For the empiricists, the trick was to show how the content of all knowledge could be reduced to pure reports of sensory experience; for the rationalists, it was showing the indubitability of the innate notions that guided and facilitated the development of knowledge. (2002, p. 125).

But the feminist critique concludes that the objectivity that is assumed within both views is unattainable by humans. And, furthermore, there is a related feminist criticism of the unrealistic enlightenment ideal of a 'transcendent disembodied subject' and an acknowledgment of the reality of 'embodied, and differentiated subjects' (Longino, 1999, p. 335). Both these criticisms parallel similar criticisms articulated within other research programs examining human rationality, including the transition from the ideal of the 'unbounded rational agent' to the 'bounded rational agent' (Gigerenzer & Selten 2002), and equivalent transitions within the 'heuristics and biases' and 'ecological rationality' research programs (Kahneman 2011, Gigerenzer 2000). Thus, research emerging from a number of disciplinary perspectives supports the view that humans are inherently biased, or to put it a different way, humans are inherently embodied and differentiated.

A central claim of the feminist critique of epistemology is that the purported 'objectivity' of certain socially dominant intellectual traditions in fact serves to hide biases that are actually present (Antony 2002, p. 114). The 'heuristics and biases' research program tends to see the uncovering of human bias as a pessimistic conclusion, humans were not as rational as had been assumed (Kahneman 2011). But feminist philosophers such as Antony and Longino do not interpret this so pessimistically.⁹ In response to the actual and ineliminable biases in human thought feminists argue that it is better epistemic practice to acknowledge the inherent biases present in all human knowledge and then to have an open and honest discussion about how to respond responsibly to those biases (Longino 1999, p. 349).

⁹ Another non-pessimistic interpretation of human cognition is central to the 'ecological rationality' research program, because humans are rational given the appropriate 'ecological' context. (Gigerenzer 2000).

Antony introduces Quine into her analysis within the historical context of his opposition to logical positivism. The positivists, she explains, were involved in 'a program of "rational reconstruction" – they wanted to show, in detail, how any empirically meaningful claim could be reduced, by the successive application of semantic and logical rules, to statements purely about sensory experience.' (2002, p. 126). But Quine's criticism of the logical positivists (or logical empiricists as they were also known), involving criticism of the analytic/synthetic distinction and his endorsement of conformational holism, had significant implications:

If logic plus data don't suffice to determine how belief is modified in the face of empirical evidence, then there must be, in addition to logic and sensory evidence, *extra-empirical* principles that partially govern theory selection. The "justification" of these principles can only be pragmatic – we are warranted in using them just to the extent that they work. (2002, p. 127-128).

Antony makes a similar point about the function of extra-empirical principles that narrow down theoretical options in relation to the concept of paradigms in the work of Kuhn:

Kuhn emphasizes that one of the chief benefits a paradigm brings with it is a degree of closure about foundational issues, instilling in members of the community a principled and highly functional unwillingness to reconsider basic assumptions. (2002, p.138).

And Antony makes the same point again with respect to worldviews:

Worldviews confer some of the same cognitive benefits as paradigms, simplifying routine epistemic tasks, establishing an informal methodology of inquiry, etc., and they also offer significant social benefits, providing a common sense of reality, and fostering a functional sense of normalcy among members of the community. (2002, p. 139).

I want to emphasize a number of points about the claims made above: (1) extra-empirical principles partially govern theory choice and that the justification of these principles can only be pragmatic; (2) accepting a paradigm (or conceptual scheme) generates a highly functional unwillingness to reconsider basic assumptions, and; (3) world views share some of the cognitive benefits of paradigms including providing a common sense of reality, and a functional sense of normalcy within a community. Embracing Sellars' advocacy of a 'stereoscopic vision' of reality (deVries, 2015) balancing both the scientific and manifest images, I suggest that the lessons learnt from the feminist critique of the philosophy of science can be applied to understanding the place of religious belief in a 'epistemically naturalised' cognitive system, when the purposes of that cognitive system are wider than predicting the future in the light of past experience.

Updating Quine's web of belief

As I mentioned before discussing the feminist critique of the philosophy of a science, I take Quine's rejection of (1) 'innate ideas' and (2) the assumption that representations of physical objects are universal aspects of human cognition, to be motivated by his empiricist outlook. But recent work in cognitive science updates our current understanding of the structure of our cognitive systems, and evidence is emerging that their structure plays a greater role than Quine assumed in the construction of our webs of belief.

First, let's consider Quine's rejection of the assumption that representations of physical objects are universal aspects of human cognition (1980, p.44). Elizabeth Spelke and Katherine Kinzler (2007) disagree with this rejection, and offer us insight into this issue by contrasting two views of human nature:

Cognitive science has been dominated by two views of human nature. On one view, the human mind is a flexible and adaptable mechanism for discovering regularities in experience: a single learning system that copes with all the diversity of life. On the competing view, the human mind is a collection of special-purpose mechanisms, each shaped by evolution to perform a particular function. The first view traces back to Enlightenment thinkers such as Locke (1689) and Hume (1748) and has been invigorated more recently by cognitive psychologists and neural network theorists (e.g. Rumelhart & McClelland, 1985; Hinton, 1993). The second view was inspired by Darwin (1871) and gained prominence with the rise of evolutionary psychology (e.g. Cosmides & Tooby, 1994; Pinker, 2002). Much public discussion has focused on the diverging ways in which these views explain human behavior. (2007, p. 89).

Spelke and Kinzler claim that neither of these views is correct. Rather, they claim that humans have four 'core knowledge systems'. These systems represent 'inanimate objects and their mechanical interactions, agents and their goal-directed actions, sets and their numerical relationships of ordering, addition and subtraction, and places in the spatial layout and their geometric relationships.' (2007, p. 89).

The object system involves cohesion, continuity, and contact, and it is present in both humans and adult non-human primates. The agent system in contrast is not based on spatiotemporal factors. Agents don't need to be cohesive, continuous, or contacting in interactions, instead agents involve goal directed, efficient, contingent, reciprocal action. Spelke and Kinzler claim that 'these findings provide evidence for a core system of agent representation that is evolutionarily ancient and that persists over human development' (2007, p. 90). Paul Bloom (2007), by stressing that the object and agent systems are dissociable, has further argued that humans are 'intuitive dualists'. Other developments in cognitive science relevant to naturalised epistemology include: the move away from a focus on the ideal (unbounded) rational agent, and toward a focus on understanding bounded rationality (Simon 1982); dual systems of reasoning (Kahneman 2011; Evans and Frankish 2009); ecological rationality (Gigerenzer 2000); and evolutionary psychology (Laland and Brown 2011).

Returning our attention to Quine's web of belief, these findings and research programs suggest that the web has some 'nodes' (or less flexible zones) within it, such that the web as a whole is not as flexible (or as 'blank') as the empiricist tradition might suggest. For example, if we do have four core knowledge systems, as Spelke & Kinzler claim, then our web of belief may be predisposed to create beliefs about 'inanimate objects and their mechanical interactions, agents and their goal-directed actions, sets and their numerical relationships of ordering, addition and subtraction, and places in the spatial layout and their geometric relationships.' (2007, p. 89). And the presence of such beliefs in our web could be more stable and inflexible than Quine may have supposed.

And now, let us turn to the issue of updating our beliefs. How does Quine imagine that we change the beliefs in our web? Quine and Ullian present a set of theoretical virtues: conservatism, modesty, simplicity, generality, and refutability (1978, p. 66-79). On the Quine/Ullian account we are to assess the beliefs within our web (or assess possible beliefs to add to our web) against these virtues. To the extent that the beliefs embody

these virtues they are good features of our web (or good additions to our web), to the extent that they do not embody these virtues they are not good features of (or additions to) the web. On this account the theoretical virtues are meant to have autonomy when determining the virtue of beliefs and only the web as whole has to stand before the 'tribunal of sensory evidence'. But the developments in cognitive science and other disciplines mentioned above constrain the web in ways that Quine did not foresee.

For example, consider the 'two-systems' approach to cognition (Kahneman 2011; Evans and Frankish 2009) such that we have a *fast and frugal* intuitive reasoning system (System 1), and a *slow and costly* deliberative reasoning system (System 2). I read Quine as if he assumes all of the web of belief exists, and is periodically reviewed, within System 2. But what if parts of the web are constructed by System 1?¹⁰ It may be the case that such beliefs are intransigent beliefs that are hard to shake out of the web of belief. Or using a different theoretical vocabulary, if we have an 'agent core knowledge system', this will surely affect the status and intransigence of 'agent' beliefs (associated with, say, non-human systems) within the web. Even using Quine's own theoretical virtues, agent beliefs (again, say, about non-human systems) may be considered simpler than non-agent beliefs. And furthermore, (if Spelke and Kinzler are right, and contra the empiricist assumption of the lack of innate ideas in humans), if agent beliefs are innate, then Quine's theoretical virtue of conservatism will reinforce their ongoing presence.

By way of illustration, consider Richard Swinburne's (1991) sophisticated dualist explanatory (and metaphysical) system involving both scientific and personal explanations. His account includes scientific explanations involving objects with powers and liabilities, and personal explanations involving persons with powers and intentions. Perhaps Swinburne is building on the intuitive attractiveness of two of our core knowledge systems.¹¹ So given the deliverances of the agent core knowledge system together with the Quinean theoretical virtues of conservatism and simplicity the result may be to retain religious belief within your web of beliefs, if that is where you began your epistemic journey.¹² In reply, Quine would presumably say that webs of belief that include religious belief are not optimal when one is pursuing the interests and purposes that he choose to motive his epistemic journey, namely 'predicting future experience in the light of past experience' (1980, p.44). But I now return to the question: *what if one had different interests and purposes? Or what if one simply had wider interests and purposes?* And here again recall Sellars' advocacy 'for a *stereoscopic vision* in which the descriptive resources of the sciences are united with the language of individual and community intentions and the dualism of the manifest and scientific images is transcended.' (deVries 2015). What if one's first priority was not only to predict future experience in the light of past experience? What if your interests and purpose extended beyond these aims? Here I invite consideration of a more expansive set of purposes of one's web of belief.

¹⁰ Perhaps the 'core knowledge systems' (Spelke and Kinzler, 2007) are all 'System 1' systems.

¹¹ Consider also that 'physical objects' and 'sets' are central to Quine's ontology (1969, p.75), and these are two of the core knowledge systems identified by cognitive science.

¹² Importantly, the religious belief may be present in the web in the first place not because it has virtue with reference to the theoretical virtues, but rather because it has virtue with reference to what I characterise as 'evolutionary virtue' later in the paper. And it should be acknowledged here that it can be considered a weakness of coherentism that it is subject to an individual's epistemic starting point.

Now let us consider Quine's rejection of innate ideas. Recall Boyer's observations that religious ideas can be understood as natural because they depend upon capacities of human brains, and that they 'are considered perfectly obvious, self-evident ideas by the people who hold them' (1994, p. 3). Both these observations point to the existence of innate religious ideas and contemporary work in CSR supports this possibility. CSR encompasses a range of research programs, but one of those programs concerns identifying an evolutionary adaptation account of the origins of religious belief based on its role in furthering cooperation within one's in-group.

Consider cooperation within groups of humans. The individuals within a group will do better if they cooperate with others within that group in comparison to if they all don't cooperate as a group. But an individual will do even better if they are within a group of co-operators, but they themselves defect on the expectation to cooperate with the rest of the group. This is a version of the 'free-rider' problem. To solve the free-rider problem groups punish defectors. But punishing defectors is itself costly (you yourself may be harmed while punishing the defector), so there is a temptation to defect on the responsibility to punish. But if a group can somehow outsource the responsibility to punish defectors then that closes off the regress of the 'free-rider problem'. For example, if it is believed that there exists a god that will punish defectors, then this motivates potential defectors to co-operate, and improves co-operation within the group. This 'supernatural punishment theory' has been advanced by Johnson and Berring (2009).

Here I take seriously the idea that our cognitive systems have been shaped by natural selection, and in particular, I take seriously the idea that the *cognitive processes that lead to belief acquisition* have been shaped by natural selection. And thus, following Johnson and Berring, I suggest, cognitive systems that generate beliefs that further social cooperation have been selected for, as in the 'supernatural punishment theory'.

So, now let us imagine that one of the purposes of one's web of belief (independently from the purpose of 'prediction') is, say, 'furthering cooperation within one's in-group'. To be clear I am not suggesting that the two purposes of 'prediction' and 'furthering cooperation' are the only purpose of a web of belief. I suspect that there are many purposes of one's web of belief. And it is important to note that the purpose of 'furthering cooperation within one's in-group' may not be a purpose that an individual is aware of. This could be a purpose given to the cognitive system of an individual by evolution. None-the-less this purpose will interact with other purposes and generate a set of theoretical virtues (and these virtues could be endorsed explicitly or implicitly). Now, importantly the theoretical virtues that are endorsed (perhaps explicitly) because of their function in selecting beliefs that predict the future in the light of past experience, will certainly have a complex pattern of interaction with theoretical virtues that are endorsed (perhaps implicitly) because of their function in selecting beliefs that further cooperation within one's in group. But that is the nature of the system I am describing. However, if this is a reasonable characterisation of the situation then what emerges is the possibility that the theoretical virtues that are endorsing beliefs for the purpose of 'furthering cooperation within one's in group' may well endorse religious belief. For example, if the 'supernatural punishment theory' is accepted, then there may

well be a selective pressure to endorse beliefs about the existence of a supernatural punisher.¹³

I have been using the phrase ‘theoretical virtue’ up to this point to label all the criteria with which beliefs are assessed within one’s web of belief regardless of the purpose those beliefs serve. But the phrase ‘theoretical virtue,’ while it might be suitable as a label for a conscious process of belief assessment, it is not the best label to capture the functioning of a process not transparent to conscious introspection. A better phrase for the criteria underlying the intuitive process might be ‘evolutionary virtue’. Thus beliefs that do a good job of ‘furthering cooperation within one’s in-group’ might be judged ‘evolutionarily virtuous’, rather than ‘theoretically virtuous’. The distinction between explicitly acknowledged theoretical virtues and implicitly functioning evolutionary virtues serves to highlight the fact that individuals do not have complete conscious control over the process that determines the content of their web of belief. And, as a brief aside, the distinction between (explicitly endorsed) ‘theoretical virtues’ and (implicitly endorsed) ‘evolutionary virtues’ might offer a way of understanding the use of the term ‘bias’ in the feminist critique of the philosophy of science.

Finally, would like to draw some parallels between the claims I am making in this paper and work being done in the CSR research program. The following is not meant to be a comprehensive analysis but rather simply suggestive of how my claims might relate to the work of others. To illustrate the parallels, I will use the work of van Leeuwen (2014). And while I don’t agree with all of the suggestions made in his paper, I do think he makes some important observations. Van Leeuwen observes that we label, for example, factual beliefs, political beliefs, theoretical beliefs, and metaphysical beliefs, all as beliefs (2014, p. 706). He claims, correctly to my mind, that this approach of identifying all these as the same sort of belief is leading to theoretical confusion. In an attempt to clear up some of this confusion he advocates a process in which the ‘functional properties’ of the belief in question can be used to distinguish different types of belief. For the record, his first attempt at this process is to distinguish between ‘factual belief’ and ‘religious credence’ and this attempt constitutes the bulk of the content of the paper. But it is not the particular focus his paper that interests me, rather I am interested in his larger project of classifying beliefs with reference to their functional properties, as this has important parallels to my suggestions here (1) that the purposes of a web of belief lead to the content of that web, and (2) that identifying the purpose of a belief (or set of beliefs) will help us understand its (or their) presence in the web.

To illustrate how Leeuwen’s work can be brought into relationship with the themes of this paper, consider his distinction between ‘factual belief’ and ‘intuitive beliefs’ (2014, p.705). Factual beliefs for van Leeuwen are vulnerable to evidence. By this he means that if one entertains a factual belief and then evidence presents itself that contradicts that factual belief then that belief should be relinquished. This directly parallels the purpose of ‘predicting the future in the light of past experience’ in the Quinean web of belief. Following Leeuwen’s suggestion to label such belief based on its functional properties we could label such belief ‘predictive belief’. So what van Leeuwen would call

¹³ I am by no means suggesting that ‘furthering co-operation within one’s in-group’ as a purpose contributing to the content of one’s web of belief is sufficient to explain all of the complexity of religious belief. I am merely suggesting that it is part of a more complex story.

‘factual belief’ I would call ‘predictive belief’ and both are evidentially vulnerable. Van Leeuwen then reflects on intuitive beliefs:

Now some “beliefs” share some but not all characteristic properties of factual beliefs. Intuitive beliefs—those embedded in our folk biology, folk psychology, folk physics, etc.—seem both to be practical setting independent and to have cognitive governance, without having evidential vulnerability. As McCloskey (1983) emphasizes, our intuitive expectations about how events unfold in the physical world are difficult to modify by education and evidence. But they govern how we imagine and are operative across practical settings. We can say for intuitive beliefs, then, that natural selection substitutes [...] for evidential vulnerability. This shaping is, however, akin to evidential vulnerability, since in both cases exposure to the physical world shapes the attitude. Much more can be said on this matter, but in any case, intuitive beliefs are close cousins to factual beliefs without being identical to them. (2014, p. 705)

There are a number of interesting ideas in this passage and van Leeuwen is certainly right when he says that much more can be said on this matter! But let me highlight a few points. Firstly, note van Leeuwen’s observation, following McCloskey, that intuitive beliefs are difficult to modify. Secondly, that intuitive beliefs can be shaped by natural selection. And finally, that natural selection substitutes for evidential vulnerability. These points suggest to me that, rather than van Leeuwen’s own suggestion that religious beliefs be re-characterised as ‘religious credences’, religious beliefs can be more accurately understood as a form of intuitive belief. Consider firstly, Boyer’s observations about the naturalness of religious ideas, secondly, Johnson and Berring’s ‘supernatural punishment theory’ as an evolutionary adaptation account of religious belief and finally my suggestion that the purposes of ‘furthering cooperation within one’s in-group’ in contrast to ‘predicting the future in the light of past experience’ might help explain the place of religious belief in one’s web of belief.¹⁴

Conclusion

And so now, in conclusion let us return to reflect upon Quine’s observation that the issue of which conceptual framework he adopted was addressed with reference to ‘the epistemological point of view’ and his further observation that this point of view is ‘one among various, corresponding to one among various interests and purposes’ (1980, p. 19). Quine’s interests and purposes focused explicitly on ‘predicting future experience in the light of past experience’ (1980, p.44).

But I have suggested that a wider set of interests and purposes inform the content of an individual’s web of belief. Perhaps as van Leeuwen (2014) suggests, we should categorise the belief like content of the mind in terms of the function of those belief like contents and then identify relevant sub-categories as appropriate. Beliefs that function to predict the future might be called ‘factual beliefs’ (van Leeuwen’s term) or ‘predictive beliefs’ (my term). But I suggest that *predicting the future* is only one of many purposes of an individual’s web of belief. Those purposes may have been chosen explicitly by the individual or those purpose may have been ‘chosen’ by evolution. Perhaps *furthering*

¹⁴ The methodology, suggest by van Leeuwen, and endorsed by me here, of categorizing and redefining putative ‘beliefs’ based on their function might also be a productive way of pursuing one form of eliminative materialism in the philosophy of mind (Churchland, 1981; Stich 1983; Churchland, 1986). Thus, what van Leeuwen calls a ‘factual belief’ and I call a ‘predictive belief’ might be better understood as a ‘predictive attitude’, in that rather than believing something, one takes a predictive attitude toward the content of a proposition.

cooperation within one's in-group is a purpose 'chosen' by evolution that leads to the population of one's web of belief with beliefs that is not there by virtue of their role in prediction. If so, then the work of Johnson and Bering (2009) on the 'supernatural punishment theory' as an evolutionary adaptation goes some way to explain the presence of religious beliefs.

Undoubtedly, predicting future experience in the light of past experience is interestingly related to furthering cooperation within one's in-group, but these are distinct purposes that lead to the endorsement of a range of interestingly related beliefs. And thus, one can endorse beliefs within the web of belief that are not there by virtue of the fact that they help us predict future experience, but rather they are there because they further our wider interests and purposes, such as furthering cooperation within one's in-group.

In a related debate, the feminist critique of epistemology and philosophy of science suggests that more interests and purposes are at play in human intellectual inquiry than are explicitly acknowledged. So, using the feminist critique we have examined, we can now consider what biases (as they would call them) are in play in an individual's web of belief. And it is interesting to note that the feminist critique does not advocate any attempt to remove the biases from our epistemic processes. The feminist critique suggests that all human epistemic processes are biased, and so they advocate, not the removal of bias, but an acknowledgment of bias and then an open and honest discussion concerning the role that bias plays in epistemic inquiry.

Thus, in conclusion, I suggest that the insights of Quine within the larger project of naturalised epistemology and the feminist critiques of both epistemology and the philosophy of science are relevant to religious epistemology. If one understands that the interests and purposes served by one's epistemic endeavours (whether chosen consciously or not) are wider than simply predicting future experience, then there may be a place for religious belief within one's web of belief. My purpose here is not to list a set of theoretical (or evolutionary) virtues that would be appropriate for a web of belief understood in the context of a broader set of interests and purposes. My purpose is simply to make the point that, just as feminists have critiqued science and made explicit the biases that they claim already exist in supposedly 'impartial' systems of theory construction, people who hold a religious perspective could do the same. All systems of theory construction are biased in some way. Thus, following the lead of the feminist critique, I suggest the way forward is to acknowledge the bias and then have an open and honest discussion. A discussion that will hopefully lead to a greater understanding of the various epistemic perspectives generated by our differing interests and purposes. And such a discussion may go some way toward answering the question: *Do religious beliefs have a place within an 'epistemically naturalized' cognitive system?*

Acknowledgements

This research is an output from a project, undertaken as part of the New Insights and Directions for Religious Epistemology Project within the Faculty of Philosophy at the University of Oxford, funded by the John Templeton Foundation. The views expressed are not necessarily those of the New Insights and Directions for Religious Epistemology Project, the University of Oxford, or the John Templeton Foundation. I thank audience

members at the presentation of this paper within the New Insights and Directions for Religious Epistemology Project and three anonymous reviewers for helpful comments.

References:

- Antony, L. (2002). Quine as Feminist: The Radical Import of Naturalized Epistemology. In L. Antony & C. Witt (Eds), *A Mind of One's Own: Feminist Essays on Reason and Objectivity* (pp. 110-153). Cambridge MA, Westview Press.
- Atran, S. (2002). *In God's We Trust: the evolutionary landscape of religion*. Oxford: Oxford University Press.
- Bloom, P. (2007). Religion is natural. *Developmental Science* 10, 147-51.
- BonJour, L. (1994) Against Naturalized Epistemology, *Midwest Studies in Philosophy* 19 (1) pp. 283-300.
- Boyer P. (1994) *The Naturalness of Religious Ideas: A Cognitive Theory of Religion*. Berkeley: University of California Press.
- Barrett, J. (2004). *Why Would Anyone Believe in God?* Lanham MD: Alta Mira Press.
- Churchland, P. M., (1981), "Eliminative Materialism and the Propositional Attitudes," *Journal of Philosophy* 78: 67–90.
- Churchland, P.S., (1986), *Neurophilosophy: Toward a Unified Science of the Mind/Brain*. Cambridge, MA: MIT Press.
- Cosmides, L., & Tooby, J. (1994). Origins of domain specificity: the evolution of functional organization. In L.A. Hirschfeld & S.A. Gelman (Eds), *Mapping the mind: Domain specificity in cognition and culture* (pp. 85–116). New York: Cambridge University Press.
- Darwin, C. (1871). *The descent of man, and selection in relation to sex*. London: J. Murray.
- deVries, W. (2015) Wilfrid Sellars, In E. N. Zalta (Ed.) *The Stanford Encyclopedia of Philosophy*, <http://plato.stanford.edu/archives/fall2015/entries/sellars/>
- Evans, J. & Frankish, K. (2009). *In Two Minds: Dual Processes and Beyond*. Oxford: Oxford University Press.
- Gigerenzer, G. (2000). *Adaptive Thinking: Rationality in the Real World*. New York: Oxford University Press.
- Gigerenzer, G. & Selten, R. (2002) *Bounded Rationality*. Cambridge MA: MIT Press.
- Glanzberg, M., (2014) Truth, *The Stanford Encyclopedia of Philosophy*, Edward N. Zalta (ed.), URL = <<http://plato.stanford.edu/archives/fall2014/entries/truth/>>.

- Hinton, G.E. (1993). How neural networks learn from experience. In *Mind and brain: Readings from the Scientific American magazine* (pp. 113–124). New York: Freeman/Times Books/ Henry Holt & Co.
- Hume, D. (1748). *An enquiry concerning human understanding*. L.A. Selby-Bigge & P.H. Nidditch (Eds.), (1975 edn.). Oxford: Clarendon Press.
- Johnson, D. and Berring, J. (2009). Hand of God, Mind of Man: Punishment and Cognition in the Evolution of Cooperation. *The Believing Primate*. Eds. Jeffrey Schloss and Michael Murray. Oxford: OUP, pp. 26–43.
- Kahneman, D. (2011). *Thinking Fast and Slow*. London: Penguin.
- Kim, J., (1988). What is “Naturalized Epistemology”? *Philosophical Perspectives, Vol. 2, Epistemology*, Ridgeview Publishing Company, pp. 381- 405.
- Kitcher, P. (1992). The Naturalists Return. *The Philosophical Review*, 101 (1), 53-114.
- Kornblith, H. (1999). In Defense of Naturalized Epistemology. In J. Greco and E. Sosa (Eds), *The Blackwell Guide to Epistemology* (pp. 158-169). Oxford: Blackwell.
- Laland, K., and Brown, G. (2011). *Sense & Nonsense: evolutionary perspectives on human behavior*. Oxford: Oxford University Press.
- Locke, J. (1689). *An essay concerning human understanding*. H. Nidditch (Ed.) (1979 edn). Oxford: Clarendon Press.
- Longino, H. (1995). Gender, Politics, and the Theoretical Virtues. *Synthese* 104, 383-397.
- Longino, H. (1999). Feminist Epistemology. In J. Greco and E. Sosa, *The Blackwell Guide to Epistemology* (pp. 327-353). Oxford, Blackwell.
- McCauley, R. (2000). The Naturalness of Religion and the Unnaturalness of Science, in F. Keil, and R. Wilson (eds) *Explanation and Cognition*, Cambridge MA: MIT Press.
- McCauley, R. and Cohen, E. (2010) Cognitive Science and the Naturalness of Religion, *Philosophy Compass* 5/9, pp. 779-792.
- McCloskey, M. (1983). Intuitive physics. *Scientific American*, 248(4), 122–130.
- Pinker, S. (2002). *The blank slate: The modern denial of human nature*. New York: Viking.
- Plantinga, A. (1995). Precis of Warrant: The Current Debate and Warrant and Proper Function. *Philosophical and Phenomenological Research*, 55 (2), 393-396.
- Quine, W. V. (1969). *Ontological Relativity and other essays*. Columbia University Press.
- Quine, W. V. (1980). *From a Logical Point of View*. Cambridge MA: Harvard University Press.

- Quine, W. V. (1995). Naturalism; Or, Living within One's Means. *Dialectica* 49:251–61.
- Quine, W. V. & Ullian, J. S. (1978). *The Web of Belief*. New York: Random House.
- Rumelhart, D.E., & McClelland, J.L. (1985). Distributed memory and the representation of general and specific information. *Journal of Experimental Psychology: General*, 114, 159–188.
- Schloss, J. & Murray, M. (2009). *The Believing Primate: Scientific, Philosophical, and Theological Reflections on the Origin of Religion*. Oxford: Oxford University Press.
- Sellars, W. (1964). Philosophy and the scientific image of man. In R. Colodny (Ed.), *Frontiers of science and philosophy* (pp. 35-78). London: Allen & Unwin.
- Simon, H. (1982). *Models of Bounded Rationality*. Cambridge MA: MIT Press.
- Spelke, E. & Kinzler, K. (2007). Core Knowledge. *Developmental Science* **10**(1), 89-96.
- Stich, S. (1983) *From Folk Psychology to Cognitive Science*. Cambridge, MA: MIT Press.
- Swinburne, R. (1991). *The Existence of God*. Oxford: Clarendon Press.
- van Leeuwen, N. (2014). Religious credence is not factual belief *Cognition* 133(3), 698-715.
- Wilson, D. S. (2002). *Darwin's Cathedral: Evolution, Religion, and the Nature of Society*. Chicago: The University of Chicago Press.
- Wood, G. (2013) The Rationality of Heuristic Religious Belief. In G. W. Dawes and J. Maclaurin (Eds) *A New Science of Religion*, New York: Routledge, pp. 189-204.
- Zagzebski, L. (1999). What is Knowledge? In J. Greco and E. Sosa (Eds), *The Blackwell Guide to Epistemology* (pp. 92-116). Oxford: Blackwell.