Studies in History and Philosophy of Biological and Biomedical Sciences 48 (2014) 21–27

Contents lists available at ScienceDirect

Studies in History and Philosophy of Biological and Biomedical Sciences

journal homepage: www.elsevier.com/locate/shpsc

Biological constraints do not entail cognitive closure

Michael Vlerick*

University of Johannesburg, Department of Philosophy, B603A, PO Box 524, Auckland Park 2006, South Africa

ARTICLE INFO

Article history:
Received 26 March 2014
Received in revised form
26 June 2014
Available online

Keywords:
Cognitive closure
Epistemic boundedness
Cognitive constraints
Epistemic scope
Extended mind
Assisted mind

ABSTRACT

From the premise that our biology imposes cognitive constraints on our epistemic activities, a series of prominent authors—most notably Fodor, Chomsky and McGinn—have argued that we are cognitively closed to certain aspects and properties of the world. Cognitive constraints, they argue, entail cognitive closure. I argue that this is not the case. More precisely, I detect two unwarranted conflations at the core of arguments deriving closure from constraints. The first is a conflation of what I will refer to as 'representation' and 'object of representation'. The second confuses the cognitive scope of the assisted mind for that of the unassisted mind. Cognitive closure, I conclude, cannot be established from pointing out the (uncontroversial) existence of cognitive constraints.

1. Introduction

The view that our epistemic activities are constrained by the structure of our minds is not new. Famously, it has been formulated by the great enlightenment philosopher Immanuel Kant (1781), who postulated that our understanding of the world is mediated by a series of a priori categories. While Kant viewed those cognitive constraints as necessary and universal (any rational—i.e. thinking—being would be endowed with those categories), Darwin's (1859) theory of evolution by natural selection cast a new light on the nature of those cognitive constraints. Indeed, in the post-Darwinian era, the human mind is revealed as a natural organ shaped by selective processes to provide our ancestors with behavioural strategies that enhanced their chances of survival and reproduction. Cognitive constraints, it appears, are contingent products of a blind evolutionary process. They are in a very real sense biological constraints. In the same way that our biology imposes constraints on diet, locomotion, and lifespan, it imposes constraints on our cognition.[1]

From the premise that the mind is a natural, limited organ with a particular structure and functioning, it seems plausible that it is ill-equipped and even downright inept to deal with certain theoretical problems. Much as we can't expect the human stomach to digest any substance (it evolved in relation to a certain kind of diet), we can't expect the mind to (be able to) represent each and every property of the world. In other words, given that the human mind is constrained by its particular genetic make-up, it is tempting to conclude that it must be closed to certain aspects or properties of the external world.

A number of prominent authors—most notably, Fodor (1981, 1983), Chomsky (1975, 1980, 1988, 2000), and McGinn (1989, 1993, 1994) have argued just this. Cognitive closure or epistemic boundedness, they argue, must follow from the fact that our minds are natural organs with a particular structure and consequently

* 9A Chichester Road, Westdene 2092, South Africa. Tel.: +27 (0) 10 559 2729 (office), +27 (0) 71 293 6749 (cell).
E-mail addresses: michaelvlerick@gmail.com, mvlerick@uj.ac.za.

[1] Although 'biological constraints' typically refer to limitations on possible body types or morphologies in biological contexts, I use it to denote that cognitive constraints are biological in nature. They are indeed the result of a particular genetic make-up (I thank an anonymous reviewer for pointing this out).

http://dx.doi.org/10.1016/j.shpsc.2014.07.008
1369-8486/

don't provide us with an all-seeing God's eye perspective or unlimited representational and processing powers. The result is what Chomsky (2000: 83) dubbed 'mysteries' or aspects of the external world that cannot in principle be known to us, because they fall outside of our epistemic reach.

Proponents of cognitive closure do not make the trivial claim that, de facto, some aspects of the world are unknown to us, nor even that some aspects will remain unknown to us. They claim that, in principle, some aspects of the world must remain unknown to beings like us. Cognitive closure follows de jure from the way our cognitive apparatus is wired. In other words, they claim that some theories *necessarily* fall beyond our cognitive scope.[2]

There are two subtly different approaches to reach this conclusion. The first is the approach from the mind. It derives cognitive closure from an analysis of the nature of the human mind. Cognitive closure is said to follow from biologically imposed cognitive constraints. The second is the approach from problems. It derives cognitive closure from an analysis of the intellectual problems that human minds have sought to solve, and arguing that a satisfactory answer to at least some of those problems falls beyond our reach. The reason they fall beyond our reach is again that our biology imposes cognitive constraints: we are constrained in the conceptual tools available to us to conceptualise solutions to these problems.

The common element in both lines of argument is that biologically determined cognitive constraints entail cognitive closure. I argue against such an entailment. Cognitive closure, I will argue, does not follow from biological constraints on cognition. At most the arguments show that cognitive closure is *possible* given cognitive constraints, but not that it is entailed. More precisely, I will identify two major unwarranted conflations at the core of the prima facie plausible claim that closure follows from constraints. The first is a conflation of what I will call 'representation' and 'object of representation'. The second confuses the cognitive abilities of the assisted mind for those of the unassisted mind.

While some authors (most notably Dennett, 1991) have claimed contra Fodor, Chomsky and McGinn that cognitive closure cannot be established from armchair reasoning and that it does not follow from a naturalised perspective on the human mind, an in-depth analysis and criticism of the fallacies at the core of arguments for cognitive closure is lacking in the literature. That is the aim of this paper.

In the next section I present both arguments for cognitive closure, focussing on Fodor, Chomsky and McGinn. In Section 3, I point out the two conflations made by advocates of cognitive closure, when deriving closure from constraints. I conclude that cognitive closure is not entailed by the existence of cognitive constraints. In this regard, the claim I am defending is not that we are not cognitively closed, but that the mere existence of biological constraints on our epistemic activities does not *entail* that we are cognitively closed. Doing so, I argue against advocates of cognitive closure such as Fodor, Chomsky and McGinn that the question of cognitive closure cannot be settled by pointing at the (uncontroversial) existence of cognitive constraints imposed by our genetic make-up.

[2] The necessity is not of a logical kind, nor is it demonstrated by a transcendental argument concerning the necessary limits of our kind of cognition. Such arguments have been advanced in modern times: e.g. Nagel (1986) suggests that certain facts about a subjective point of view cannot be represented from an objective point of view, nor from other subjective points of view. This sort of argument, and the logical necessity it invokes, is not what Chomsky, Fodor and McGinn have in mind (and is not the topic of this paper). Rather, the kind of necessity at stake is physical. Beings like us—wired with the particular kind of brains we possess—are cognitively closed.

2. Two arguments for cognitive closure

2.1. The argument from the mind

According to Fodor (1983: 120) 'a psychological theory represents the mind as epistemically bounded if it is a consequence of the theory that our cognitive organization imposes *epistemically significant constraints* on the beliefs that we can entertain' (my italics). I take it that 'epistemically significant constraints' mean constraints that block certain true and non-trivial beliefs or representations of the world. If those constraints would not prevent us from producing an accurate representation of some relevant aspect of the world, they would not be consequential or epistemically significant. Fodor's notion of epistemic boundedness, in this regard, aligns with the notion of cognitive closure as defined above. This epistemic predicament, according to Fodor, is the fate of all empirical (i.e. naturalised) theories of mind, since any such theory allots some kind of endogenous structure to the mind. Therefore, he argues, epistemic boundedness or cognitive closure (i.e. the claim that some true beliefs about the world fall outside our conceptual reach) must follow. In Fodor's words:

> It then seems to me hard to see how the unboundedness view can be made empirically plausible. The point is that any psychology must attribute some endogenous structure to the mind (really unstructured objects – bricks, say – don't have beliefs and desires and they don't learn things). And it is hard to see how, in the course of making such attributions of endogenous structure, the theory could fail to imply some constraints on the class of beliefs that the mind can entertain (Fodor, 1983: 125).

The argument can be formalised as follows:

F1: Human minds are endogenously structured minds.
F2: Endogenously structured minds are cognitively constrained minds (weak entailment—'hard to see how, in the course of making such attributions of endogenous structure, the theory could fail to imply some constraints on the class of beliefs that the mind can entertain').
F3: Cognitively constrained minds are epistemically bounded/cognitively closed minds (strong entailment—Fodor construes these constraints as constraints on epistemically significant beliefs, entailing that some non-trivial beliefs about the world are unthinkable).
F4: Human minds are epistemically bounded/cognitively closed minds.

This argument is valid, but I will argue that one of its premises is false. More precisely, I will reject F3.

A similar argument for cognitive closure comes from Chomsky (2000). Chomsky conjectures that our capacity for science is due to a biologically endowed 'science forming faculty'. This 'science forming faculty' or SFF is loosely defined as comprising those aspects of the human mind that 'enter into naturalistic inquiry' (Chomsky, 2000: 83). While this faculty enables us to engage in science, it also constrains it. Some aspects of the world lie beyond the reach of this faculty and remain necessarily 'mysteries' to the human mind. We are closed off in principle from ever knowing them. Note the close similarity with Fodor's argument. From the premise that human cognition or science is grounded in an innately determined (and therefore fixed) cognitive structure or faculty (in Chomsky's case a 'science forming faculty') (F1) both authors derive the existence of constraints on human thinking or science (F2), and from this premise (the existence of constraints) both authors derive closure (F3). Fodor postulates 'epistemically significant constraints' and Chomsky 'mysteries'.

Indeed, the concept of cognitive closure or epistemic boundedness is at the core of Chomsky's (2000: 83) problem–mystery distinction. Problems are questions we could (in principle) answer or aspects of the world we could explain. Mysteries, on the other hand, are insoluble or inexplicable in principle, given our biological make-up. The mechanism of heredity, for instance, was an unsolved problem before Crick and Watson's discovery of DNA, not a mystery. However, according to Chomsky's line of argument, some true theories about the world can be expected to fall outside the scope of the human science forming faculty. Whatever those theories would bring to light, therefore, are mysteries to the human mind. Why? Because 'like other biological systems SFF [science forming faculty] has its potential scope and limits' (Chomsky, 2000: 83).

There is a striking analogy with Chomsky's (1965) well-known view on language and language acquisition. Much as our language 'module' determines a Universal Grammar endowing human speakers with access to only a subset of all theoretically possible grammars, our science forming faculty determines a set of concepts and principles endowing human thinkers with access to only a subset of all theoretically possible (successful) theories about the world. Non-human speakers could have a different Universal grammar, speaking languages which could not be acquired by us—at least not in the way young children acquire a first language despite a so-called poverty of stimulus.[3] Similarly, non-human thinkers—'alien scientists'—could have a different science forming faculty, giving rise to theories about the world which would necessarily remain beyond our ken.

Inferring epistemic boundedness from looking at the mind, Fodor and Chomsky do not single out those areas in which we are bounded. There is indeed no way to pinpoint mysteries. How is one going to show that a true representation of the world cannot be conceptualised by a human mind? Fodor (1983: 123) seems sensible to this point, (merely) arguing that he doesn't see how 'any remotely plausible cognitive theory' could conceivably guarantee epistemic unboundedness. Similarly, Chomsky (2000: 83) doesn't think a sharp demarcation between problems and mysteries can be drawn. We can only assume that some of the unsolved problems are in fact mysteries. McGinn, however, takes a stronger stance. His approach is to single out problems which—he argues—can never be answered by the human mind.

2.2. The argument from problems

Rather than looking at the mind and inferring that some successful theories must fall outside its scope, McGinn (1993) looks at problems and infers that the mind is not equipped (in principle) to come up with a solution. He advocates a 'Transcendental Naturalism'. The human mind is cognitively closed to the solution of certain problems, according to McGinn, not because those problems are different in nature than solvable scientific problems, but because the particular structure of our minds obstructs that knowledge. This is what he means by 'Transcendental Naturalism': there are problems that transcend our cognitive capacities, but at the same time are not supernatural in their solutions; they are not ontologically different from the natural problems we can solve. Examples of such problems are: the mind-body problem, the self, meaning and intentionality, free will, a priori knowledge, and knowledge in general—the problems which have typically raised philosophical perplexity throughout the history of human thought. McGinn expands on the mind-body problem:

> How is it possible for conscious states to depend upon brain states? How can technicolour phenomenology arise from soggy grey matter? What makes the bodily organ we call the brain so radically different from other bodily organs, say the kidneys – the body parts without a trace of consciousness? How could the aggregation of millions of individually insentient neurons generate subjective awareness? We know that brains are the de facto causal basis of consciousness, but we have, it seems, no understanding whatever of how this can be so. It strikes us as miraculous, eerie, even faintly comic. (McGinn, 1989: 349).

McGinn (1989: 350) suggests that the mystery arises because we are 'cut off by our very cognitive constitution from achieving a conception of that natural property of the brain (or of consciousness) that accounts for the psychophysical link'. We do not, in other words, have the conceptual tools at our disposal to understand (and therefore represent) the nexus between mind and body, or between consciousness and brain. This problem, that has been haunting philosophers for ages, cannot therefore be solved by a human mind—not because of the nature of the problem (it is, according to McGinn's Transcendental Naturalism, intrinsically no different from problems we do solve) but because of the nature of our cognitive apparatus.

In this regard, McGinn's argument is an inference to the best explanation. The best explanation for the fact that we can't come up with a solution to certain problems is a limitation of the mind (rather than an ontological difference between problems we can solve and problems we can't). To support this claim, McGinn turns to the working of the human mind. As he points out: 'ideally, TN [Transcendental Naturalism] needs to be accompanied by a worked-out theory of human cognitive capacity, from which it would be demonstrable that certain forms of understanding are not humanly accessible, or run against the cognitive grain' (1993: 18). While McGinn admits that formulating a full-fledged theory of how the mind works is not realistic, he nevertheless claims that 'some stab should be made at saying what at least such a theory would look like' (18). The result of this stab is the so-called 'CALM' conjecture, which stands for 'Combinatorial Atomism with Lawlike Mappings'. Our mode of thought, according to McGinn, is essentially combinatorial. The way we understand a domain is by identifying primitive elements subject to 'specified principles of combination which generate determinate relations between complexes of those elements'. In other words, we grasp the world and its phenomena by accounting for it in terms of a set of primitive elements and their 'lawlike' interaction.

Once again we can formalise McGinn's argument:

M1: All humanly graspable representations of the world are CALM-like
M2: Some aspects of the world are not representable in CALM-like fashion
M3: Some aspects of the world are not humanly graspable

I will reject M1. McGinn bases M1 by looking at some influential sciences. According to McGinn, physics, linguistics and mathematics 'clearly exemplify the pattern' (19). Indeed, physics postulates primitive elements and the 'resulting complexes (macroscopic material objects) are then governed by lawlike relations [between these primitive elements]'. Similarly, in linguistics we have primitive elements (words, phonemes) which are combined by combinatorial rules (grammar) to form complex wholes. Finally,

[3] According to Chomsky (1965) children pick up the grammatical structure of the natural language they acquire despite the fact that they are not presented with sufficient linguistic data to identify the grammatical rules and patterns by mere inductive inference. Therefore, Chomsky argues, given that language knowledge is underdetermined by the input available for learning, we must be endowed with an innate theory of language—a so-called 'Universal Grammar'.

mathematics—especially geometry according to McGinn—has its primitives (lines, planes and volumes) which are then combined by means of a set of rules into ever more complex structures.[4]

Philosophical perplexity in the face of problems such as the relation between mind and body, the nature of the self, meaning and intentionality, and free will, arises from the fact that our CALM mode of thought is utterly inept to grapple with these problems. 'Conscious states' McGinn claims, 'are not CALM-construable products of brain components' (1993: 37), much as the self and free-will, among those other problems, resist being answered in terms of the particular framework our minds must employ in thinking about the world.

Pinker (1997) who is sympathetic to McGinn's claim, points out something similar. Studying the working of the mind, he argues, 'we can glimpse *why* certain problems are beyond our ken' (564, his italics). The mind, he points out, owes its power to its syntactic, compositional, combinatorial abilities. The typical philosophical problems, however, have a peculiar 'holistic and everywhere-at-once and nowhere-at-all and all-at-the-same-time' character. Consciousness can't be constructed out of the sum of brain states, the self not out of our various body parts and/or brain states and free will is by definition not a causal chain of events and states. These enigmas, Pinker ponders, may arise from 'a mismatch between the very nature of these problems and the computational apparatus that natural selection has fitted us with' (Pinker, 1997: 564–565).

3. Two fallacies

3.1. Conflating representation and object of representation

Let us assume with Fodor and Chomsky that there are constraints on human thinking (F2). In other words, the human mind cannot entertain any theoretically possible representation of the world. The fact that our biology imposes constraints on what the human mind can produce is hardly controversial. Indeed, it is a particular natural organ with its particular modus operandi. Alien scientists, possessing radically different minds, could very well come up with ways of representing the world that are not and could never be intelligible to the human mind. From the premise that we cannot conceive of all theoretically possible successful representations about the world, it is tempting to conclude that there must be aspects of the world that we cannot represent (i.e. aspects to which we are cognitively closed) (F3). This, I oppose to Fodor and Chomsky, does not follow. The existence of cognitive constraints does not entail that those constraints are epistemically significant (block true and non-trivial beliefs about the world) or that there must be mysteries. Cognitive constraints do not entail cognitive closure.

A distinction has to be made between representation and object of representation. The distinction is subtle but crucial. 'Representation'[5] refers to the way we access the world and the 'object of representation' to the entity or property in the world which is represented (given that we are dealing with a representation which successfully refers). Language provides us with a helpful analogy. The utterance 'the cat is on the mat' comprises a set of phonemes which given the English vocabulary and grammar conveys a certain meaning. The representation is that very utterance for English speakers. The object, however, is the actual cat on the mat. An 'object' which is also represented by the French expression 'le chat est sur le paillasson' or the Dutch equivalent 'de kat is op de mat'.

The fact that we only have a subset of all theoretically possible representations at our disposal does not entail that we only have access to a subset of objects or properties in the world. Different kinds of representations can represent the same thing, and being closed to represent a thing in a particular way, doesn't necessarily entail that one is closed to representing that thing tout court. Take, for instance, two creatures, each with a particular ability to represent a point in space. The first represents a point in space by using numerical coordinates on an axis system, defining a particular location in terms of the intersection of a numerical value on a horizontal axis and one on a vertical axis. The second, on the other hand, zooms in on a particular point by means of a grid system. It represents space in terms of frames, each subdivided in further frames. A particular point in space will then be represented in terms of the particular 'sub-sub-sub… frame' that corresponds with it. This second creature could very well be closed to the mode of representation of the first. The fact, however, that it has a different kind of representation of spatial locations at its disposal does not entail that it cannot represent the very same thing as the first, i.e. the same point in space to the same degree of precision. In a similar vein, Nagel's (1974) bat representing the location of a fly through echo-location, represents it differently than the frog's visual representation. Both however succeed in representing its location.

The fact that the set of representations at our disposal are constrained by the way our sensory and cognitive apparatuses are wired and are therefore both particular (the result of one evolutionary path among many) and contingent (other types of cognisers, 'alien scientists' could very well access the world with different kinds of representations), does not mean we can (in principle) only represent a particular subset of parts and properties of the world. Similarly, the fact that some theoretical possible representations could very well be unintelligible to the human mind, does not entail that the kinds of representations available to us, at best, represent the world in a limited way. Again the analogy with language is helpful. Assuming Chomsky (1965) is correct about our possession of a universal grammar, there are certain constraints to the structure of (all) possible human (natural) languages. This, however, does not entail that certain contents are necessarily inexpressible, merely that the way in which they must be expressed is constrained by our nature.

What could motivate the conflation of representation and object of representation and consequently the entailment of cognitive closure from representational constraints, is what Goldman (1986) dubbed 'the mirror metaphor of correspondence'. The mirror metaphor holds that every state of the world determines a singular corresponding representation. For every state of the world, it claims, there is only one accurate representation, and the properties of the corresponding representation are—therefore—solely determined by the state of the world it represents. This often implicitly held 'mirror metaphor' is central to the traditional doctrine of correspondence. An influential version of the mirror metaphor of the doctrine of correspondence is, according to Goldman (1986: 151), developed in Wittgenstein's Tractatus (1928). The 'Tractarian version of correspondence' states that the world is a totality of facts and that a proposition is true if and only if it corresponds with a fact. A true proposition, in this regard, is an objective depiction or a true reflection of an external state of affairs.

[4] Incidentally, McGinn's CALM structure does not seem to fit all sciences. Statistic driven sciences, such as economics, climate sciences or epidemiology, do not appear to employ this basic framework. While this already undermines M1, I will offer principled reasons why M1 fallacious (cf. 3.2).

[5] I am not committing to any theory of intentionality. My aim is not to deal with the philosophical questions of how representation is possible or what representation is, but merely to distinguish cognitive access from the object in the external world that is being accessed. In doing so, I take a successful representation to be a representation which stands for an object/property in the world for a particular cognitive being. How—let alone if—this is possible and what this relation consists of, is not of my concern here.

The mirror metaphor, however, is only one possible metaphor for correspondence between a representation and its object. Goldman (1986: 152) substitutes the mirror metaphor of correspondence for what he considers to be the preferable metaphor of 'fittingness'. 'Fittingness', he explains, in the 'sense in which clothes fit a body'. This metaphor allows for the 'categorizing and statement-creating activity of the cognizer-speaker', while—at the same time—'capturing the basic realist intuition that what makes a proposition or statement true is the way the world is' (152). The subject does, in other words, capture the properties and elements of the world by means of its own (particular and contingent) cognising activity. Goldman expands on the theme:

> There are indefinitely many sorts of apparel that might be designed for the human body, just as there are indefinitely many categories, principles of classification, and propositional forms that might be used to describe the world. […] Despite all this variety there is still the question, for any specified type of apparel, whether a specific token of that type fits a particular customer's body. The question of fittingness is not just a question of style of garment. It depends specifically on that customer's body. Similarly, although the forms of mental and linguistic representation are human products, not products of the world per se, whether any given sentence, thought sign, or proposition is true depends on something extra-human, namely, the actual world itself (Goldman, 1986: 152–153).

The fact, in this regard, that the same entity or property in the world can be represented in more than a single way, does not entail that a particular way cannot be right or wrong. There remains, indeed, a connection between a representation and the object which is represented: the representation is or should be 'moulded' to fit its object. Much as we can use clay, stone or bronze to bring out the same shape in matter, there are various possible ways to represent the external world. While the world itself does not need to be represented in any particular way, any particular representation needs to fit the world in order to be accurate. Different representations can latch onto a single property or part of the world and being endowed with a mere subset of all possible modes of representation does not entail closure.

Therefore, if one is willing to let go the ultimately misguided mirror metaphor, the approach from the mind fails. More precisely (F3) is unwarranted. Cognitive closure does not follow from cognitive constraints. Deriving closure from constraints stems from an unwarranted conflation of representation and the object of representation; a conflation of the way we access the world with the actual parts and properties we access. It strikes me as very unlikely, however, that Fodor, Chomsky and McGinn would adhere to the mirror metaphor and a literal correspondence theory of truth. Indeed, assuming that the mind is a natural organ with its contingent structure and functioning (F1 and M1), how could they make sense of such a literal conception of correspondence between our cognitive products and the entities they are meant to represent? Given that our cognitive abilities and by extension our representations are the outcome of a contingent evolutionary path, it seems odd to claim that they objectively mirror the things they represent (implying that the only way in which the world can be accurately represented is the way in which our evolved faculties happen to represent it). It makes much more sense from a natural mind perspective, I think, to claim that our cognitive apparatus endows us with a particular, species-specific way to make sense of and represent its surrounding. A particular way among other possible ways.

Moreover, reasoning from a literal conception of correspondence, what grounds would Fodor, Chomsky and McGinn have to claim that even the least mysterious aspects of the world correspond to our (contingent) cognitive outputs. What guarantee would McGinn have that some aspects of the world are intrinsically CALM-like? Similarly, what guarantee would Fodor have that some properties and elements of the external world can be perfectly mirrored by the set of concepts we have at our disposal? As Goldman (1986: 155) points out, we have no reason to assume that the world 'contains factlike entities' or is structured as our language and thought. From a literal correspondence criterion of truth, no knowledge is warranted. At least not if we abandon the Aristotelian view of an unconstrained mind endowed with universal rationality. Given that proponents of cognitive closure do grant us epistemic access to some parts and properties of the world, we can assume that they are committed to a looser criterion of correspondence and therefore jump to the conclusion of cognitive closure from merely establishing constraints on representing the world.

3.2. Conflating assisted mind and unassisted mind

The potential human epistemic reach should not be gauged by looking at the unassisted mind and senses. The question of cognitive closure, therefore, cannot be settled by looking at our 'naked' biological endowment. Indeed, human beings have radically extended the scope they have on the world, provided by their bare senses and unassisted mind. We detect distant galaxies with powerful telescopes and calculate their distance with complex mathematical equations. Supported by these cognitive levers, we even managed to overthrow substantial parts of our intuitive understanding of the world (e.g. folk physics and its conjecture that every object 'strives' towards rest, or folk biology and its view that species are endowed with an immutable essence) and replace them with scientific theories (in casu, Einsteinian physics and Darwinian evolutionary biology). An answer to the question of cognitive closure, therefore, should take the assisted or 'extended' mind and senses into account, not the biological, unassisted sensory and cognitive faculties.

According to Clark and Chalmers (1998) and Sterelny (2010), science and its impressive cognitive feats are primarily the product of our ability to extend our minds' capacity through interacting with its environment. Rather than to our internal computational engine, we owe our best theories of the world to what Sterelny (2010) calls 'scaffolds'. Scaffolds are the external resources we use to enhance our cognitive reach.[6]

First, science is not done by isolated individuals. There is a division of cognitive labour. The production of (scientific) knowledge is not the product of a single unassisted mind. Einstein did not come up with relativity from scratch. Every scientific theory is the result of cooperation, both on a horizontal level (groups of cooperating scientists rather than isolated individuals) and a vertical level (cumulative progress made over many generations).[7]

[6] I am not committing to any particular metaphor for assisted cognition. Whether the mind is 'extended'—as proposed by Clark & Chalmers (1998), suggesting that mental processes can be (spatially) extended beyond the brain—or 'scaffolded', suggesting that cognitive processes are enhanced by external structures without therefore being spatially extended, is not of my concern here. The point I am making (and which both Clark and Chalmers and Sterelny endorse) is simply that external resources can induce cognitive augmentation.

[7] I am not committing to any kind of theory about the development of science. Whether it is continuous (to some extent) as for instance Boyd (1989), Psillos (2009) and Worrall (1989) have claimed or characterised by discontinuity and distrupting revolutions (Feyerabend, 1975; Kuhn, 1962), is not of my concern. I merely point at the uncontroversial observation that scientists do not go about their activity from zero, but build to an important extent on what previous generations have developed (if only by making use of some core aspects of science such as mathematics, the empirical method, etc).

Tomasello (2001), in this regard, argues that our ability to read the minds of conspecifics is the key cognitive adaptation accounting for human culture (and by extension science). This faculty allows us 'to pool their cognitive resources both contemporaneously and over historical time in ways that are unique in the animal kingdom' (2001, 135). In a similar vein, Boyd, Richerson, and Henrich (2011) argue that, rather than to innate intelligence, we owe our success to "our uniquely developed ability to learn from others" (1).

Second, science is not done by 'naked' minds. It is supported by a series of cognitive artifacts such as mathematics—which provides us with a powerful system to represent and compute quantitative information. Furthermore various notation devices—i.e. writing—enable us to store huge amounts of information and work out complex mathematical equations (De Cruz, 2008). The leverage of these scaffolds can hardly be overestimated. They radically transform our ability to retain and process data, enabling us to produce theories about the world which would otherwise remain unthinkable. Imagine having to reach an Einsteinian or even Newtonian worldview without mathematics.

Finally, we have a myriad of technological instruments at our disposal, enabling us to detect data we could never gather by means of our naked senses (e.g. telescopes, microscopes, barometers, and the like), and process information in ways we could never do with our unassisted mind (computers). Moreover, future technological innovations can be expected to radically enhance the scope and depth with which are able to investigate the universe.

It should be clear that our cognitive scope on the world—the data we detect and the representations we form based on these data—is not determined by the scope provided by our unassisted biological endowment. Even stronger, what makes us smart and what sets us apart from non-human animals—as a number of authors have pointed out (Boyd et al., 2011; Carey & Spelke, 1994; Tomasello, 2001; Vlerick, 2012)—is not so much our innate intelligence but how we manage to leverage it. Much as we shouldn't gauge the lifting power of a man operated crane, by the physical strength of the operator, we shouldn't gauge the limits of human knowledge by assessing the limits of its innate cognitive faculties. This, however, is exactly what proponents of the cognitive closure thesis do.

We are epistemically bounded, according to Fodor, because the class of concepts we can entertain is endogenously constrained (F3). In other words, the concepts we can produce to represent the world are determined by our innate cognitive wiring. Note that Fodor defends the counterintuitive thesis known as 'concept nativism', in which every primitive concept at our disposal—i.e. a concept that cannot be decomposed into other concepts—(including, for instance, the concept of 'proton') is innate. Given that we have but a particular set of innately determined primitive concepts at our disposal and no means to form new concepts except by recombining primitive concepts into complex ones, Fodor concludes that our minds must be bounded. Indeed, they are bounded to an innately determined 'potential conceptual endowment'. As Fodor (1981: 314) puts it 'we have no guarantee that the concepts required to build true science are situated in that space'.

Apart from the fact that Fodor's radical concept nativism seems very implausible at the outset (did Aristotle's innate stock of concepts include the conceptual building blocks to refer to protons and computers?), it fails to account—as Rellihan (2005) points out—for the deferential acquisition of concepts. Indeed, Rellihan argues, acquiring a concept does not need to entail the capacity to entertain the constituents of the concept (237). We can very well acquire a concept—i.e. determine its extension—by deferring to experts or scientific instruments. A lay botanist, for instance, can acquire the concept 'Elm' by including all those trees which have been tokened 'Elm' by an expert. Similarly, we can acquire the concept 'acid' by deferring to a scientific instrument—in this case a litmus paper—and token a solution as acidic whenever the paper turns red. A telling example Rellihan brings to our attention is that of a colour blind man acquiring the concept of red by tokening all those objects red which his wife points out as red. While he does not possess the qualia red, he is nevertheless able to discriminate all red objects from non-red ones. Doing so, he augments his potential conceptual endowment, which given his visual impediment did not include colours. While I disagree with Rellihan's claim that in these examples new concepts are acquired, to the extent that concepts are representations not 'things represented', he does point at something important. Indeed, in the context of the distinction between 'representation' and 'object of representation' introduced in the previous section, the use of external elements such as other people and instruments, enables us to access and therefore represent things in the world for which our unassisted sensory and/or cognitive faculties do not have a ready-made representational media. The colour blind man can pick out the things (red objects) without having the qualia, much as the scientist can pick out the property (acidity) without having the proper sensory receptors to detect it. In short, access to things in the world is not determined by the access provided by our unassisted biological endowment. Cognitive closure therefore, which is precisely the claim that access to certain parts and properties of the world is in principle denied to us, should not be gauged by looking at our 'naked' biological endowment.

Chomsky, on the other hand, argues for cognitive closure based on our possession of an innate 'science forming faculty', with its potential scope and limits. Our cognitive reach, according to Chomsky, can be determined by looking at this innate faculty, much as our linguistic reach (the subset of logically possible languages which are intelligible to human speakers) is determined by our language module. Again, our epistemic limits are assessed by looking at the mind's innate endowment. However, as pointed out, what makes us smart or what extends our cognitive reach is not so much our strict innate endowment but the way we are able to leverage it.

Finally, McGinn derives cognitive closure from a mismatch between the nature of certain problems and the wiring of our mind (cf. the CALM conjecture). Once again, the focus in on the naked, unassisted mind. In a telling passage, McGinn (1994: 140) exclaims that we cannot fly. 'Thus we learn to walk quite naturally and everyone is pretty efficient in this department; our swimming abilities, however, are laboriously acquired and show much individual variation; and when it comes to flying, well, it shouldn't even be attempted'. Similarly, he argues, 'we can survey our cognitive skills to see where we are strong, weak and downright inept'. What McGinn fails to acknowledge is that while we cannot flap our arms and take off, we are actually capable of flying… in airplanes. We might be unable to fly with our naked, unassisted bodies, but we sure succeed in flying with the proper support. The same goes for cognition. With our bare mind, we couldn't even make sense of some of the simplest problems which are easily solved by a reasonably mathematically competent 10-year old armed with pencil and paper. Equating our epistemic capabilities to the power of our unassisted mind is seriously misguided.

To this, McGinn might reply that when dealing with these allegedly 'unsolvable' problems, no amount of 'scaffolding' will help, since our lack of understanding stems from a qualitative issue, not a mere lack of quantitative processing power. No computer—however powerful—will magically yield the solution to the mind-body problem. But cognitive scaffolds do not merely boost our reasoning power. They radically alter our reasoning modes. Consider how far science has distanced itself from our commonsense

understanding of the world. In the pre-Einsteinian era one could very well have argued—and in fact the Kant did argue—that the human mind could only conceive of time and space as independent, absolute entities. This would mean that relativity theory would be humanly unthinkable. This goes to show that we are not bound to even our most tenacious intuitive representations and ways of reasoning. Indeed, supported on a series of powerful scaffolds such as mathematics, telescopes and the like, Einstein's thinking was able to override this intuitive representation of time and space which is firmly anchored in our cognitive nature. In other words, the assisted human thinker (Einstein in this case) was able to reject a perspective on the world any unassisted human thinker necessarily holds, and replace it by a different perspective. Assuming that Einstein's theory is an (approximately) accurate representation of certain properties of the universe, one could say in retrospect that the unassisted human cogniser is cognitively closed to those properties. This, however, can obviously not be said for the assisted thinker.

In this regard, given that cognitive scaffolds can radically alter our conceptual abilities, McGinn's blunt assertion that the mind is bound to CALM-like reasoning (M1) implies that we have a clear view on which external cognitive tools we have at our disposal and how this could scaffold our reasoning, and more importantly, which tools we could potentially have at our disposal. This, I take it nobody will deny, is plainly wrong. Much as Aristotle could not conceive of the epistemic tools that led to Einstein's relativity theory, there is absolutely no way of predicting what tools might be generated in the future and what they might bring to light. Cognitive closure, therefore, does not follow from an assertion of problems which are unsolved at present and a hasty 'stab' at how the mind works. (M1) is—apart from being a poor fit to all human epistemic endeavours (see footnote 2)—seriously misguided. To prove cognitive closure, rather than focussing on the unassisted mind and its biological constraints, one needs to show how no possible scaffolding could off-set cognitive closure in certain areas given the cognitive constraints determined by our biology. The above mentioned authors have certainly not done this and therefore have failed to establish cognitive closure.

4. Conclusion

Cognitive closure, defined as a principled inability to represent certain parts and properties of the world, does not follow in principle from biologically determined cognitive constraints. I argue against Fodor, Chomsky and McGinn, that we cannot simply infer closure from the existence of such constraints on our epistemic activities. I do admit however that, given these constraints, representing some aspects and properties of the external world might turn out to be beyond our cognitive reach. What I reject, in this regard, is not the possibility of cognitive closure, but its necessary entailment from the existence of biological constraints on cognition. In this regard, I argue that the question of cognitive closure cannot be settled by pointing out constraints and epistemic pessimism with regards to stubborn conundrums is both premature and unwarranted.

References

Boyd, R. (1989). What realism implies and what it does not. *Dialectica, 43*, 5–29.
Boyd, R., Richerson, P., & Henrich, J. (2011). The cultural niche: Why social learning is essential for human adaptation. *Proceedings of the National Academy of Sciences of the United States of America, 108*(2), 10918–10925.
Carey, S., & Spelke, E. (1994). Domain-specific knowledge and conceptual change. In L. Hirschfeld, & S. Gelman (Eds.), *Mapping the mind: Domain specificity in cognition and culture* (pp. 169–199). Cambridge: Cambridge University Press.
Chomsky, N. (1965). *Aspects of the theory of syntax*. Cambridge: MIT Press.
Chomsky, N. (1975). *Reflections on language*. New York: Pantheon.
Chomsky, N. (1980). *Rules and representations*. New York: Columbia University Press.
Chomsky, N. (1988). *Language and problems of knowledge: The Managua lectures*. Cambridge: MIT Press.
Chomsky, N. (2000). *New horizons in the study of language and mind*. Cambridge: Cambridge University Press.
Clark, A., & Chalmers, D. (1998). The extended mind. *Analysis, 58*(1), 7–19.
Darwin, C. (1859). *On the origin of species by means of natural selection, or, the preservation of favoured races in the struggle for life*. London: John Murray.
De Cruz, H. (2008). An extended mind perspective on natural number representation. *Philosophical Psychology, 21*(4), 475–490.
Dennett, D. (1991). *Review of McGinn, the problem of consciousness*. The Times Literary Supplement (May 10, 1991)
Feyerabend, P. (1975). *Against method: Outline of an anarchist theory of knowledge*. London: New Left Books.
Fodor, J. (1981). The present status of the innateness controversy. In J. Fodor (Ed.), *Representations* (pp. 257–316). Cambridge: MIT Press.
Fodor, J. (1983). *The modularity of mind: An essay on faculty psychology*. Cambridge: MIT Press.
Goldman, A. (1986). *Epistemology and cognition*. Cambridge: Harvard University Press.
Kant, I. (1781). Critique of pure reason. In P. Guyer, & A. Wood (Eds.), *The Cambridge edition of the works of Immanuel Kant*. Cambridge: Cambridge University Press (1998)
Kuhn, T. (1962). *The structure of scientific revolutions*. Chicago: University of Chicago Press.
McGinn, C. (1989). Can we solve the mind-body problem? *Mind, 98*, 349–366.
McGinn, C. (1993). *Problems in philosophy: The limits of inquiry*. Oxford: Blackwell.
McGinn, C. (1994). The problem of philosophy. *Philosophical Studies, 76*(2), 133–156.
Nagel, T. (1974). What is it like to be a bat? *The Philosophical Review, 83*(4), 435–450.
Nagel, T. (1986). *The view from nowhere*. Oxford: Oxford University Press.
Pinker, S. (1997). *How the mind works*. New York: W.W. Norton and Company.
Psillos, S. (2009). *Knowing the structure of nature. Essays on realism and explanation*. London: Palgrave MacMillan.
Rellihan, M. (2005). Epistemic boundedness and the universality of thought. *Philosophical Studies, 125*, 219–250.
Sterelny, K. (2010). Minds: Extended or scaffolded. *Phenomology and the Cognitive Sciences, 9*, 465–481.
Tomasello, M. (2001). Cultural transmission: A view from chimpanzees and human infants. *Journal of Cross-Cultural Psychology, 32*, 135–146.
Vlerick, M. (2012). How can human beings transgress their biologically based views? *South-African Journal of Philosophy, 31*(4), 717–735.
Wittgenstein, L. (1928/2001). *Tractatus logico-philosophicus*. New York: Routledge.
Worrall, J. (1989). Structural realism: The best of both worlds? *Dialectica, 43*(1–2), 99–124.