Darwin's doubt - Implications of the theory of evolution for human knowledge

Michael Marie Patricia Lucien Hilda Vlerick

Proefskrif ingelewer vir die graad Doktor in die Wijsbegeerte in die Fakulteit Lettere en Sosiale Wetenskappe aan die Universiteit van Stellenbosch.

Dissertation presented for the degree of Doctor of Philosophy in the Faculty of Arts and Social Sciences at Stellenbosch University.

Supervisor: Dr. J.P. Smit December 2012

Verklaring

Deur hierdie proefskrif elektronies in te lewer, verklaar ek dat die geheel van die werk hierin vervat, my eie, oorspronklike werk is, dat ek die alleenouteur daarvan is (behalwe in die mate uitdruklik anders aangedui), dat reproduksie en publikasie daarvan deur die Universiteit van Stellenbosch nie derdepartyregte sal skend nie en dat ek dit nie vantevore, in die geheel of gedeeltelik, ter verkryging van enige kwalifikasie aangebied het nie.

Declaration

By submitting this dissertation electronically, I declare that the entirety of the work contained therein is my own, original work, that I am the sole author thereof (save to the extent explicitly otherwise stated), that reproduction and publication thereof by Stellenbosch University will not infringe any third party rights and that I have not previously in its entirety or in part submitted it for obtaining any qualification.

June 8th 2012

Abstract

In this dissertation I enquire into the status, scope and limits of human knowledge, given the fact that our perceptual and cognitive faculties are the product of evolution by natural selection. I argue that the commonsense representations these faculties provide us with yield a particular, species-specific scope on the world that does not 'correspond' in any straightforward way to the external world. We are, however, not bound by these commonsense representations. This particular, species-specific view of the world can be transgressed. Nevertheless, our transgressing representations remain confined to the conceptual space defined by the combinatorial possibilities of the various representational tools we possess. Furthermore, the way in which we fit representations to the external world is by means of our biologically determined epistemic orientation. Based on the fact that we are endowed with a particular set of perceptual and cognitive resources and are guided by a particular epistemic orientation, I conclude that we have a particular cognitive relation to the world. Therefore, an accurate representation for us is a particular fit (our epistemic orientation) with particular means (our perceptual and cognitive resources).

Abstrak

Hierdie tesis handel oor die aard, omvang en limiete van kennis, gegewe dat ons perseptuele en kognitiewe vermoëns die resultaat van evolusie deur middel van natuurlike seleksie is. Eerstens, word daar geargumenteer dat die algemene voorstellings wat hierdie vermoëns aan ons bied 'n partikuliere, spesie-spesifieke siening van die wêreld aan ons gee, wat nie op 'n eenvoudige manier korrespondeer aan die werklikheid nie. Ons is egter nie gebonde aan hierdie voorstellings nie. Hierdie partikuliere, spesie-spesifieke siening van die wêreld kan oorskry word. Ons is egter wel beperk tot die konseptuele ruimte wat gedefinieër word deur die kombinatoriese moontlikhede van die voorstellingsmiddele tot ons beskikking. Verder word die manier waarop ons hierdie voorstellings aan die wêreld laat pas deur ons biologies gedetermineerde epistemiese oriëntasie bepaal. Dus, gegewe dat ons 'n spesifieke stel perseptuele en kognitiewe vermoëns het en deur 'n spesifieke kognitiewe epistemiese oriëntasie gelei word, staan ons in 'n spesifieke kognitiewe verhouding tot die wêreld. 'n Akkurate voorstelling (m.a.w. kennis vir ons) is om spesifieke vermoëns (perseptuele en kognitiewe vermoëns) op 'n spesifieke manier (epsitemiese oriëntasie) aan die wêreld te laat pas.

Acknowledgements

First and foremost, I would like to thank Dr. J.P. Smit (Stellenbosch University), my supervisor, for guiding me through this endeavour. This dissertation is the product of the countless discussions we had, as well as the valuable advice, and the careful review on his part.

Furthermore, I would like to express my gratitude to Dr. Chris Buskes (Radboud University Nijmegen) for his interest in this project, the numerous exchanges of ideas, and for proofreading substantial parts of this thesis.

Finally, my thoughts go to my parents for their trust and support, and, of course, to my wife, Séverine, for standing by my side every step of the way, and for leaving everything behind to embark with me on this adventure.

Stellenbosch, June 8th 2012

Contents

Abstract iii

Acknowledgements v

Introduction 1

Chapter 1. A particular intuitive view of the world 5

1. Introduction 5
2. Perception 6

2.1 Vision 6

2.1.1 Physical cause

2.1.2 Vision as problem-solving

A) Judging distances

B) Discriminating objects

C) Reliable illusions

2.1.3 Conclusion

2.2 Auditory perception 11

2.2.1 Physical cause

2.2.2 Hearing as problem-solving

A) Reconstructing the origin of sound

B) Clearing up the input

2.2.3 Conclusion

2.3 Other senses 15

2.4 Conclusion 16

2.4.1 Perceptual closure

2.4.2 The contingent causal relation between stimuli and percept

2.4.3 The constructive role of the sensory apparatus

3. Intuitive theories 19

3.1 Innate predisposition to individuate subject matters 19

3.2 Folk 'sciences' 21

3.2.1 Folk physics

3.2.2 Folk biology

3.3 Conclusion 26

3.3.1 Innateness

3.3.2 Scientific inaccuracy

3.3.3 Species-specificity

4. Why do we perceive and make sense of the world the way we do? 28

4.1 Perception 29

4.1.1 Detecting 'affordances'

4.1.2 Using the perceptual input

4.1.3 Human perception

4.1.4 Umwelt

4.2 Cognition 35

4.2.1 Increasing control over the environment

4.2.2 Adaptive problems

4.2.3 Particular intuitive assumptions about the world

5. Conclusion 39

Chapter 2. The Lorenzian fallacy: deducing truth from functionality 41

1. Introduction 41

2. Lorenz's evolutionary epistemology 42

2.1 Ontogenetic a priori as a phylogenetic a posteriori 43

2.1.1 Adaptations as representations

2.1.2 A non-arbitrary a priori

2.1.3 The experience of the lineage

2.1.4 The enigmatic congruence between mind and the world

2.2 Hypothetical realism 48

3. Uncovering the fallacy 50

3.1 Plantinga's case against evolutionary naturalism 50

3.1.1 Possible relations between beliefs and behaviour

3.1.2 Belief-cum-desire systems

3.1.3 Conclusion

3.2 Natural selection and rationality 54

3.2.1 Lorenzian evolutionary epistemology's problematic assumptions

3.2.2 Does evolution produce optimally well designed systems?

3.2.3 Are optimally well designed systems rational ones?

3.2.4 Cognitive biases and error management theory

3.2.5 Non-cognitive purposes

3.2.6 Conclusion

3.3 The mechanism of evolution 61

3.3.1 Undirected variation

3.3.2 Evolution as a gradual, non-teleological process

3.3.3 Path dependency

3.3.4 The environment

A) Vollmer's mesocosm

B) A stone age mind

3.3.5 Genetic drift and genetic hitch-hiking

3.3.6 Conclusion

3.4 The lack of external perspective 70

3.4.1 How does the mind mirror the world?

3.4.2 Analogy with other organisms

4. Conclusion 73

Chapter 3: Transgressing our biological bias 75

1. Introduction 75

2. The unique human ability to transgress its biological bias 76

3. What enables us to 'transgress'? 79

3.1 The blank slate model 80

3.2 The Jamesian model 83

3.3 The operational model 86

4. Theories about what makes us smart 87

4.1 Mapping across domains 87

4.2 Analogy and metaphor 89

4.3 Explicit and conscious representations 91

4.4 Metarepresentational thought 92

5. Necessary conditions for transgressing 94

5.1 E.T.s on an icy planet 95

5.2 What cognitive faculties would E.T. need to transgress its commonsense view? 96

5.3 Framework of necessary cognitive faculties for transgressing 98

6. Integrating human cognitive faculties in the framework 100

6.1 Representing the representation and its parts 100

6.2 Variation through recombination 102

6.3 Epistemic value system 103

7. Conclusion 106

Chapter 4: Scope and limits of transgression 108

1. Introduction 108

2. Scope of transgression 109

2.1 Productivity of thought 109

2.2 Representational flexibility 110

2.3 Open cognitive relation 111

2.4 Conceptual spaces: reconciling openness and closure 113
3. Limits of transgression 116
3.1 Bias from representing the world in a particular way 116
3.1.1 Sensory input
3.1.2 Intuitive grasp
3.2 Closure from representing the world in a particular way 124
3.2.1 Sensory resources
3.2.2 Cognitive resources
4. Conclusion 135

Chapter 5: Implications for human knowledge 137

1. Introduction 137
2. What is the relation between our representations and the external world? 138
2.1 Saving epistemological realism: the distinction between grasp and content 138
2.2 The nature of cognitive grasps: impossibility of a God's eye view 143
2.3 The fallacy of correspondence as 'mirroring' 146
2.4 The nature of the fit 148
3. What are the consequences of our epistemic relation? 152
3.1 The two-sided determination of knowledge 152
3.2 The threat of relativism 154
4. What are the limits to our epistemic relation to the world? 159
4.1 Limits set by fallibility 159
4.2 Limits set by computational capacity 162
4.3 Limits set by scope 165
5. Conclusion 171

Chapter 6: Evolutionary epistemology's double-edged sword 173

1. Introduction 173
2. Evolutionary epistemology's double threat 173
3. Answering the double threat 175
    3.1 Is my argument self-defeating? 175
    3.2 Is my argument circular? 178
4. Defence of a naturalised epistemology 179
    4.1 No alternative 180
    4.2 Virtuous circularity 182
5. Conclusion 184

Conclusion 186

References 190

Introduction

In a letter to William Graham, Charles Darwin (1881) expressed the following concern:

> With me the horrid doubt arises whether the convictions of man's mind, which has been developed from the mind of the lower animals, are of any value or at all trustworthy. Would anyone trust in the convictions of a monkey's mind, if there are any convictions in such a mind?

The theory of evolution by natural selection – developed by Darwin (1859) and Wallace (1858) – has, indeed, consequences that extend far beyond the science of biology. The philosopher Daniel Dennett (1995) likens it to a 'universal acid'. 'Darwin's idea', he argues, 'eats through just about every traditional concept, and leaves in its wake a revolutionized world-view, with most of the old landmarks still recognizable, but transformed in fundamental ways' (63). Epistemology is no exception. It is affected to its very core by this corrosive idea, since – as Darwin (1881) himself fully realised – the theory of evolution sheds a whole new light on the origin and therefore the scope and limits of the human mind.

Age-old epistemological questions, in this regard, have recently been recast in the light of the theory of evolution. Such evolutionary approaches to epistemology have led to two distinct research programs. The first, which Bradie (1986) labels the 'evolution of epistemic mechanisms', reasons about human knowledge from the premise that our cognitive (and perceptual) faculties are the product of evolution by natural selection. The second, the 'evolution of epistemic theories' (Bradie, 1986), on the other hand, is concerned with the evolution of ideas or theories themselves, using models and metaphors drawn from biological evolution. This dissertation forms part of the former research program, gauging the status, scope and limits of human knowledge, from the premise that our perceptual and cognitive faculties are the product of evolution.

Typically, evolutionary considerations on human knowledge – in the sense of Bradie's (1986) 'evolution of epistemic mechanisms' – have given rise to two opposite positions. The first argues that natural selection will shape our cognitive faculties in such a way that they produce representations of the world which correspond (to a large extent) to the external world, since accurate representations will enhance an organism's chances of survival and successful reproduction and – therefore – the faculties shaping those representations must have been selected. The second – in contrast – argues that we cannot expect our cognitive faculties to produce accurate representations of the world, since natural selection will *only* endow us with faculties generating representations which lead to survival and reproduction enhancing behaviour and that there are other, non-adaptive evolutionary forces at work.

In this dissertation, I will reject both positions. I reject the former 'evolutionary justification arguments', claiming that natural selection ensures the correspondence of our representations to the world, based on the counter-arguments provided by the opposite sceptical camp. I do, however, also reject the sceptical conclusion this camp reaches, based on what I'll refer to as the distinctive human ability to transgress the biologically based, commonsense representations it holds of the world. We are, indeed, able to overcome the biologically determined representations evolution endowed us with (i.e. we are able to represent the world in ways that go beyond and against these uncritical representations) and are therefore, not bound by these particular representations, which cannot – as rightfully pointed out by the sceptics – be expected to correspond to the external world.

The status, scope and limits of human knowledge, therefore, have to be considered in the light of this distinctively human cognitive ability to transgress the uncritical representations it holds in virtue of its perceptual and cognitive nature, shaped by the process of evolution by natural selection. Homo sapiens is, indeed, as opposed to any other species on this planet, free to conceptualise the world in an unlimited number of contingent ways. This remarkable cognitive ability is the true hallmark of the human epistemic situation. An analysis of our ability to transgress will, consequently, enable us to shed new light on those epistemological questions.

Human knowledge, I will argue, is both constrained and free. It is constrained in the sense that it is the product of a particular and contingent set of perceptual and cognitive resources. It is free in the sense that it is not restricted to the biologically based representations that evolution endowed us with. This insight is at the core of this dissertation. Radically thinking through the consequences of what I'll call our 'open epistemic relation to the world within the boundaries of a particular conceptual space', I will gauge both the status of knowledge (what makes a representation true?) as well as its scope and limits.

Plan of the thesis

Chapter 1 looks at the way we are biologically predisposed to view the world, as the result of the perceptual input we gather from it through our senses and the innate cognitive knowledge systems we apply to this input. I will argue that evolution provided us with a particular species-specific scope on the world that does not 'correspond' in any straightforward way to the external world. Our senses do, indeed, gather but part of the potential stimuli from the world, and the stimuli to which they react give rise to a particular phenomenal 'percept'. Our cognitive predispositions, on the other hand, interpret this input in the light of intuitive theories which are at odds with the modern sciences.

Chapter 2 looks at the argument of Konrad Lorenz (1941, 1973) which states that, since our perceptual and mental abilities are the product of natural selection, the representation of the world they provide us with must be (approximately) accurate. I will reject this claim, arguing that we cannot expect evolution to provide us with accurate representations. This reinforces the conclusion that the uncritical view of the world we hold in virtue of our senses and cognitive apparatus is a contingent, species-specific view, which cannot be expected to correspond (even approximately) to the external world.

In Chapter 3, however, I will point out that this particular, species-specific view of the world, as the result of the perceptual and cognitive apparatus that the blind process of evolution provided us with, can nevertheless be *transgressed* – i.e. that those commonsense representations can be substituted by different and often contradictory representations we perceive as epistemically preferable. I will argue that the possibility of transgressing our biologically based views is grounded in the three-fold cognitive ability to metarepresent, to produce alternative representations with the available resources and, when doing so, to be guided by an epistemic orientation.

Chapter 4 will gauge the scope and limits of these transgressing representations. I will argue that we can both shape an infinite amount of possible representations and that we are not bound to represent any subject matter in a particular way. Therefore, I will characterise the cognitive relation we entertain with the world as an 'open cognitive relation'. This, however, does not imply that there are no limits to the representations the human mind can produce. Those limits, I will argue, are set by the particular senses and cognitive reasoning patterns that evolution provided us with. They either bias us against or straight-forwardly close us off to certain possible representations of the world. Our epistemic situation, therefore, while providing us with an open epistemic relation to the world, is – nevertheless – comprised within a particular conceptual space.

Chapter 5 will bring the main argument of this thesis to a close, looking at the implications this 'open cognitive relation within the boundaries of a particular conceptual space' holds for our epistemic endeavours. In doing so, it will consider the issue of epistemological realism, analysing whether our representations can correspond to the external world and – if so – what the nature of this correspondence is. Furthermore, it will look into the threat of relativism and outline possible sources of limitation to a successful epistemic relation with the world.

Chapter 6, finally, defends the approach taken in this dissertation against the two most basic threats evolutionary approaches to epistemology face. In this regard, I will argue that my argument resists both the threat of being self-defeating and the threat of being circular.

Chapter 1: A particular intuitive view of the world

1. Introduction

In this chapter, I will look at our intuitive view of the world - the way the world appears to us on a commonsensical level. Reasoning that this view is the product of both the perceptual input we receive through the senses and our innate predisposition to frame this perceptual input in intuitive assumptions about the world, I will conclude that we are endowed with a particular, species-specific understanding. This intuitive understanding of the world can be opposed to a theoretical or critical understanding, which attempts to overcome our commonsensical grasp by identifying and questioning the assumptions and methods of inference underlying it. The possibility of 'transgressing' our intuitions about the world will be the subject of chapter 3.

As pointed out above, my claim is that we perceive the world in a *particular* way and hold *particular* intuitive assumptions about the world, yielding a *particular* view of the world. 'Particular', in the sense I use it, can be defined as: a contingent way of viewing (perceiving and understanding) the world, both not necessary and not universal. This implies, on the one hand, that the properties of the world do not force us to perceive and understand them in the way that we do (not necessary) and, on the other hand, that other organisms could view the world differently (not universal).

Regarding perception, I will argue that we are perceptually closed to certain elements of the world, that there is a causal but contingent relation between stimuli and percept, and that our sensory apparatuses have an active role in creating the percept, i.e. that they add to the content of the percept. Therefore, I conclude that what is given in perception does not mirror the world, but is a *particular* representation of it.

Regarding our intuitive assumptions about the world, I will look at our innate predisposition to 'carve up the world' into different categories and to apply intuitive theories to each of them. These underlying, commonsensical assumptions about the structures of the world, I will argue, can be shown to be innate, scientifically inaccurate and thoroughly species-specific, meaning that they are tailored to the specific needs of our species. They are, in other words, *particular*: both not necessary – the world can be understood differently, as achieved by the framework of modern science – and not universal – they provide us with a species-specific framework.

Both how we perceive and how we interpret this perception is, indeed, the outcome of an evolutionary process that selects *particular* features in the light of *particular* biological needs in a *particular* environment. These abilities evolved in order to increase our control over our environment, guiding our actions in ways that enhance our chances for survival and reproduction, not to give us an accurate, objective and complete understanding of reality. Therefore, I conclude, we are endowed with a particular intuitive view of the world.

2. Perception

In this section, I will first give a scientific overview of the working of our senses. This includes the physical causes triggering the different senses and the way our perceptual abilities process these data to provide us with useful information (cf. vision and hearing as problem solving). Based on this scientific account, I will then draw general conclusions regarding the relation between the external world and our perception of it.

2.1 Vision

2.1.1 Physical cause

Vision works by the projection of light onto the retina. A light source emits photon particles that move in a more or less straight line at, of course, the speed of light. Photons are units of light energy. The flight path of a photon is called a ray. Our visual experience, in this context, is made up by the distribution and directionality of photon flow entering the pupil and projecting onto the retina (Boynton, 1979:47-48). By doing so, the flow stimulates photoreceptors (rods and cones), which in turn pass a neural signal to the brain (Pinker 1997:215).

Before entering the eye, the photons are reflected off surfaces of objects. When they hit a surface, depending on the reflectance of the surface, a certain percentage of the light cast upon the surface will be reflected. Depending on the pigmentation and the texture of a surface, it will absorb a certain percentage of light. The more a surface absorbs light – or photon particles – the fewer photon particles will bounce off. An object that mostly absorbs light, therefore, will appear black, while one that mostly reflects light will appear white (Levine & Schefner, 1991:326).

What we see, therefore, are light particles or photons reflecting (or not) on all kinds of surfaces, each time reflected in different amounts depending on the properties of that very surface. That is what makes us aware of the presence of objects and their particular colour. Colour is the visual experience we derive from the wavelength of photons. The human eye is sensitive to a very narrow portion of the electromagnetic spectrum. It perceives wavelengths in the range of 375 to 750 nm (nanometer). The shortest perceptible wavelengths are perceived as violet or reddish blue. As the wavelength increases, it is perceived as blue, then as green, then as yellow, and finally as red, which is the way we perceive the longest wavelengths of the spectrum (Levine & Schefner, 1991:387-388; Boynton, 1979:48) .

Different wavelengths of light, caused by reflection off different surfaces, are perceived as different colours. Natural light contains almost all wavelengths. The light emitted by the sun contains all wavelengths in approximately equal amounts. This light appears white to the human observer. White, therefore, is the least 'pure' colour – it contains all wavelengths in equal amounts, not one particular part of the visible spectrum (Levine & Schefner, 1991:388).

2.1.2 Vision as problem-solving

How can we turn a myriad of light and colour into a useful, clear picture in which objects stand out and distances can be judged? In order to do this, our visual apparatus has to turn two-dimensional (reversed) projections on the retina – caused by billions of photons reflecting off surfaces – into three-dimensional mental hypotheses about the spatial layout of the perceived environment. There is an important (re)construction at work here, considering that the retinal projection holds no information whatsoever about the third dimension. The two main aspects to be reconstructed from this two-dimensional projection are the distance of particular objects and the discrimination of objects against their background (Pinker, 1997:215). Furthermore, the brain needs to deform and obstruct much of the original input in order to provide us with a useful, intelligible representation (Granit, 1977: 128).

A) Judging distance

Every single point on the retinal image can be caused by a point at any possible distance from the eye (Pinker, 1997:215). Without this constructive ability to hypothesise about the relative depth of every perceived point, both from one another and from the position from which one is looking, we would see nothing but moving colours crammed together. This is not an easy task to complete, and yet, a very crucial one. Without depth-vision, we wouldn't gather much useful information about our surrounding environment, perceiving only a kaleidoscopic myriad of colours. How do we solve this problem and obtain three-dimensional information about the distance of objects based on a two-dimensional projection onto our retina?

We owe this ability to the mechanism of stereoscopic vision. Each eye has a slightly different view. This is called 'binocular parallax'. These two pictures have to be united into a single picture. Every point, in this regard, is in a slightly different position on each retina. This very problem, however is the source of the solution. Indeed, the distance of every particular point can now be inferred from the difference in position that this point occupies in the projection onto each retina. Based on the angle formed by the eyes and their separation in the skull, the relative difference of every point in the projection in both eyes can now be instantly 'calculated' to infer the distance of the source of the perceived point (Pinker, 1997:219-222).

This system, however, is not infallible. The brain has to detect the same mark in both views and unite them. This matching problem is responsible for visual illusions caused by repeated patterns (often on wallpaper), where looking at it you often see one of the patterns leaping out, creating false perspective. This happens when two different patterns are connected together as if they were one and the same point, seen from a slightly different angle by both eyes, and therefore at some distance (Pinker, 1997:225).

B) Discriminating objects

The other important problem to solve in order to obtain a three-dimensional visual experience based on a two-dimensional projection is perspective. How do we see three-dimensional shapes by drawing upon a two-dimensional projection of colours onto our retina? We appear to have what Pinker (1997:243) calls a 'shape analyzer'. It infers from the retinal image what the most probable state of the world is. In order to do so, it is equipped with both an innate theory of projection – how do objects appear in projection? – and an innate theory about the world – what kind of objects does it have (244)? This, however, does not suffice by itself. We complete this picture with our sense of lightness and colour (245). Different kinds of matter – as we have pointed out above – reflecting back different wavelengths, give us the perception of different shades of colour and brightness. The problem is that shades of colour and brightness also depend on the level of illumination. So in order to deduce an object's material, our 'lightness analyzer' must try to factor out the level of illumination (246). This is done by making further assumptions. The first one is that the lighting is uniform – in other words, that the whole scene is either in the sun, in the shade or in the dark. Different levels of lightness, therefore, are the result of the different matter of objects on which the light reflects. The second assumption is that the world is a rich mixture of wavelengths (i.e. different colours). The final assumption is that gradual changes in brightness and colour are the result of illumination, while abrupt changes are caused by boundaries of objects (247). Yet another problem to solve in order to discriminate between objects and to see their shapes is what Pinker refers to as 'the effect of slant on shading'(248). This problem arises as a result of the fact that a surface facing a light source will reflect back a lot of light, while a surface angled parallel to the source reflects much less. The same amount of reflected light could therefore be reflected from a darker surface facing the light or from a lighter surface angled away. Once again, we fall back on another assumption, namely that the surface of the object uniformly reflects back light. Our 'shape-from-shading analyzer', in this light, is fooled by the moon.

Since it is pockmarked with craters, it does not reflect back light uniformly and, therefore, looks like a disc rather than a sphere to us (249).

This is obviously a fallible way of deducing what really is the case. However, the different analysers – making hypotheses about shape, lightness and shade – are all taken into consideration when deducing a three-dimensional picture from a two-dimensional projection. Together, they provide us with a hypothetical and fallible representation, but one that works fine in most cases (Pinker, 1997:249).

C) Reliable illusions

As shown above, in order to provide us with a useful visual representation, our visual apparatus reconstructs a three-dimensional image, based on a two-dimensional projection, using a variety of cues and innate expectations. We also perceive different wavelengths and amounts of light as different shades of colour and brightness, to yield an orderly and useful representation of the surroundings. However, this does not suffice. As Ragnar Granit (1977), Nobel laureate for his work in visual physiology, points out, the brain must deform much of its informational content in order to make our visual experience intelligible.

> We must not underestimate what the interpreting brain itself adds to make the seen world more intelligible than does a pure peripheral input, dependant though the cortex is on information from feature detectors. The purposive brain requires a considerable degree of invariance, size constancy, a fixed verticality, approximately invariant surface colours, some constancy of velocity and direction of movement and, above all, a steady world; in short, a large number of what one is fully entitled to call 'reliable illusions'. They are all constant errors with respect to the informational content of the primary sensory message. (Granit, 1977:128).

Vision, in this light, does not reflect the environment as accurately as possible, but provides us with a clear, simplified picture. In order to achieve this, our visual apparatus constantly deforms the original sensory input, yielding a mental picture stripped of its chaotic complexity.

2.1.3 Conclusion

The problem that visual perception faces becomes clear at this point – how to deliver a useful representation of the world, based on an enormous amount of potential stimuli? First of all, in order not to ‘cram’ our vision with useless data, we only respond to a small portion of the electromagnetic spectrum. This relevant part is then experienced as different shades of colour and brightness, to make us perceive objects against a neutral background of natural lighting. This coloured projection is immediately processed by different innate modules in our brain to deliver a full-fledged – albeit hypothetical – three-dimensional image of our surroundings. This is how we get an intelligible representation of our surroundings. If, on the other hand, we merely received the two-dimensional projection on our retina as input, vision would provide us with nothing more than an indistinct blur. It is therefore precisely because of its selective receptiveness, its particular way of experiencing stimuli, and its ability to make hypotheses about a three-dimensional layout, that vision provides us with useful information.

2.2 Auditory perception

2.2.1 Physical cause

What do we hear? What is the physical nature of sound? Similarly to light, sound consists of waves. However, those waves are not a type of electromagnetic radiation – as light is – but a purely mechanical phenomenon. Sound consists of changes in air pressure, generated either by vibrations of objects (e.g. by knocking on something) or by a release of air (e.g. whistling or speaking). This change in air pressure then propagates as a wave, moving in all directions as ever increasing circles around the source. The speed at which those waves travel depends on the medium through which they travel. Sound travels at approximately 340 m/sec through air, considerably faster through water, and even faster through metals. The frequency of a sound, measured in Hertz (Hz), is the number of times a particular waveform is repeated per second. The period, on the other hand, is the amount of time that one particular waveform lasts (Warren, 1999:1).

Just as our visual apparatus is only sensitive to a particular portion of the electromagnetic spectrum, our auditory apparatus only perceives frequencies ranging from more or less 20 Hz to 20 kHz. The frequency of the sound wave is related to the perceived pitch of the sound in hearing, just as the wavelength of photon particles are perceived as colour. The amplitude of the sound wave – the amount of pressure or energy – is related to the perceived loudness of the sound. However, other factors also come into play - frequency, for instance, also influences perceived loudness (Levine & Schefner, 1991:476-477, 485).

When sound waves reach our ear, they are processed through three different stages. First, they reach the outer ear - the visible part of the ear - which – as I will explain later – contributes to the localisation of the source of the sound. Reflections in our pina - our outer ear - enacts an acoustic transformation and leads the pressure changes through our ear canal, ending at the eardrum. This ear canal does more than just pass on the sound. It works as a powerful amplifier, comparable to a resonant tube. Those amplified pressure changes then cause the eardrum to vibrate. This vibration is then picked up and transmitted by a chain of three small bones, or ossicles, located in the middle ear. At this point the air-borne pressure waves are converted into liquid-borne waves. Normally this would mean a loss of 99.9% of the power of the wave. However, three mechanical levers, amplifying the sound each time, enable the hearing system to make the transition without too much loss. The inner ear, finally, contains the receptors responsible for hearing, converting those pressure waves into experienced sound, linking its frequency to perceived pitch and its amplitude to perceived loudness (Warren, 1999:5-12).

2.2.2 Hearing as problem-solving

A) Reconstructing the origin of sound

The main problem that hearing must solve is to reconstruct the origin of the sound. Without this information, hearing wouldn't be a very useful sense, leaving us with an indistinct brouhaha. With respect to this problem, however, things are not as simple as they appear. Unlike light waves, sound waves do not travel in a straight line. Rather, they behave as expanding circles, much like the circles of waves on water after dropping a stone on the surface. This makes it very difficult to locate their source. In order to accomplish this feat, we appear to be using different cues involving both ears together (binaural) or each ear separately (monaural) (Warren, 1999:29).

A first binaural cue is the difference in intensity. When the left ear is directly in the path of the sound waves, and the other is blocked by the head, an inference is made about the spatial origin of the sound. There are two ways by which the sound can then reach the right ear. It can either bend around the head or pass through the skull. Low-frequency sounds (long wavelengths) can easily bend around the head and, in doing so, will reach the right ear without being significantly blocked. High-frequency sounds, on the other hand, cannot bend around the head and will have to go through the head. In doing so, the head will cast an 'auditory shadow' over the right ear. It will filter and reduce the amount of stimulus that reaches the right ear. Therefore, the sound reaching each ear will differ slightly in intensity, leaving a cue as to where the sound originates from. (Levine & Shefner, 1991:506). Another binaural cue is the slight difference in time at which each ear – being at a different distance from the source – will receive the sound stimulus (507).

These cues alone, however, could not provide us with a satisfying result. But we have another trick up our sleeve. We can move our head, changing the stimuli and comparing. This sensitively increases the information we gather from these cues (508). In addition to the cues involving both ears, we also have monaural cues. As pointed out above, the outer ear is equipped with a characteristic shape that significantly helps us to locate the source of the sound. Batteau (1964) investigated this and concluded that the corrugation of the pinnae (outer ears) produces echoes and has an effect on the intensity of high-frequency components. These echo-induced intensity transformations provide us with information concerning the direction as well as the elevation of the sound source.

But direction isn't all that matters – estimating the distance to the sound source also offers valuable information. The first and most obvious cue is the decrease in intensity correlated with the increase in distance. The intensity of a perceived sound is approximately inversely proportional to the square of the distance of its source (to the listener). This means that a twofold change in distance corresponds more or less with a fourfold change in intensity.

However, this still does not enable us to distinguish a loud sound from afar from a faint one close by. Another cue, however, adds to this information – the ratio of direct to reverberant sound. The more reverberant components of the sound (echoed components) reach the listener in relation to direct components, the further off the source is thought to be (Warren, 1999:51-52). A last cue, finally, works for far off sounds only. It appears that high frequencies carry less far than low ones. Thunder, for instance, is perceived as a deeper, lower rumble from far away than when it is close by, since only those frequencies cross the distance. Therefore, the lower the sound, the further off its origin is thought to be (54).

B) Clearing up the input

The mechanics of hearing and the ways of utilising several cues to further determine direction and distance to the source of noise, however, do not suffice to provide us with the clear audible information we gather from our environment. Similarly to vision, in order to obtain a useful, clear 'picture' of what's going on, a thorough selection and (re)construction has to be performed by the brain, cancelling out irrelevant sounds that often obstruct important signals. Warren points out that we need to pick up relevant sounds out of a swamp of (often more intense) insignificant noise. If we could only hear the loudest sounds, hearing would lose much of its usefulness. Therefore, the auditory system is endowed with mechanisms giving us access to the fainter sounds. Furthermore, we even seem able, under some circumstances at least, to restore sounds that have been obliterated (Warren, 1999:134).

Signals of importance can, of course, be completely masked and therefore unperceivable. But when signals are only partially masked, leaving audible snatches of the signal before and after the obliterated segment, we are able to reconstruct these parts. This reconstruction happens unconsciously, leaving the listener to believe that what he hears is an unobstructed, perfectly clear signal – as if there were no missing segments at all. Both components of familiar nonverbal sounds and missing components in speech can be restored. In the latter, of course, linguistic skills also come to the rescue in inducing from the context which word or fragment of a word is appropriate (Warren, 1999:135-136). The importance of restoring important missing sounds is enormous. If we weren't endowed with this faculty, much of the relevant information gathered through our auditory faculty would be lost in the loud and meaningless brouhaha of the environment.

2.2.3 Conclusion

Just as is the case for vision, in order to provide us with useful information, hearing is selective, triggered by only a small part of the range of frequencies, giving us a particular phenomenal experience derived from stimuli – perceiving frequency as pitch – and making hypotheses about the state of the world – the direction and distance of the origin of the sound – by using several cues. Furthermore, much of the irrelevant input is obstructed and much of the lost input is reconstructed, yielding a clear perception.

2.3 Other senses

Both smell and taste, the so-called 'chemical senses', are triggered by particular substances. These substances, when coming in contact with taste-buds or olfactory receptors, cause a particular phenomenal quality, the taste or smell. Other substances do not trigger the receptors and remain unperceived. (Levine and Shefner, 1991:573-574, 592-593).

Somato-sensory sensation, on the other hand, is caused by stimuli directly in contact with our body. Four submodalities can be distinguished within somato-sensory sensation, although there is some overlap. The first one is 'proprioception' – literally: perception of oneself – the awareness of the position of the body and the limbs in space. The second one is 'tactile sensation' – the non-painful stimuli sensed when something is placed against the body surface. The third one is 'nociception' or pain - the sense elicited by noxious stimuli. The final one is 'temperature' – the sense elicited by stimuli either warmer or colder than the body surface (Levine and Shefner, 1991: 545). Interestingly, the feeling of pain is triggered by separate receptors, not merely by an excess of stimuli exerted on the touch and temperature receptors (when one suffers from an excessive pressure on the body or extreme temperatures in contact with the skin). Some receptors are sensitive to noxious mechanical stimuli, others to thermal stimuli and some respond to both. Pain, therefore, has to be viewed as a completely independent modality, evolved to give us additional information (551).

In this regard, different realms of phenomenal qualities are derived from stimuli directly in contact with the body surface. Similarly as for vision and audition, where wavelengths are perceived as colour and frequency as pitch, certain properties of the somato-sensory stimuli give rise to particular perceptual experiences. We feel hot and cold, different kinds of pain, soft, rough, tickly, etc., because of some particular physical properties of the substances in contact with our body.

2.4 Conclusion

2.4.1 Perceptual closure

Our senses are triggered by only a small part of the available stimuli. First of all, they respond only to stimuli within a small range. Our visual receptors are sensitive to a narrow portion of the electromagnetic spectrum, our auditory receptors to frequencies above 20 Hz and below 20 kHz, our chemical senses only to particular substances, and our somato-sensory receptors to particular properties of elements in the environment (e.g. temperature, pressure exerted by objects on our skin, texture of surfaces, etc.). Other organisms perceive different elements of the environment, because the range to which their sensory apparatuses are sensitive differs. Bees, for instance, are known to see UV-light and dogs hear frequencies above 20 kHz.

Secondly, the stimuli within the perceivable range need to be strong enough for us to perceive them. Our senses only provide us with a certain level of resolution. We do not see Mars or even the craters on the moon with the naked eye, because they are too far off, nor do we perceive sounds that are too faint, and the same goes for smells, tastes and tactile sensations, for that matter. Other organisms have senses providing them with more detailed levels of resolution. Eagles, for instance, are known to see more sharply than we do, dogs smell better, and so on.

Finally, there are realms of potential information for which we have not evolved appropriate receptors at all and which, therefore, remain unnoticed by us. Some migrating birds, for instance, are endowed with ‘magnetoception’, the perception of the earth’s magnetic fields, providing them with information about direction and latitude. Those magnetic fields, however, remain irremediably outside our perceptual radar, because we are not endowed with the proper sense-organs to be receptive to those stimuli.

Therefore, we cannot escape the conclusion that we perceive but part of the world. Our senses are only triggered by a small range of potential stimuli that, moreover, need to be strong enough. Furthermore, some stimuli are not perceived, simply because we do not possess the necessary sensory abilities to be triggered by them. In this context, each organism has a particular window on the world. Whatever falls outside of that window remains irremediably hidden to the organism.

2.4.2 The contingent causal relation between stimuli and percept

The account of the physical cause of the senses tells us that wavelengths of rays of photon-molecules are perceived as colour, changes in air pressure are perceived as sound, pressure and temperature of objects in contact with our skin are processed by different kind of receptors to provide us with a typical phenomenal quality, and finally, some chemical properties are detected by olfactory and gustatory receptors, making us experience a particular smell and taste.

The picture we gather from this is that the causal originators of perception – i.e. the external stimuli – are mediated by the perceptual organ to deliver the percept – i.e. the phenomenal experience. This mediation is the result of the working of our particular sense organs. This implies that the same stimulus could, in principle, give rise to a different percept, given a different mediation. In other words, different organisms with different sense organs could experience the same stimuli in different ways. Indeed, nothing about the physical causes of perception, whether they are reflecting light molecules or suddenly displaced air molecules or an interaction of external molecules with some of our own – as in the case of somato-sensory and chemical perception – forces us to perceive them the way we do. While changes in perception correspond to changes in the world, the detection of the latter could have been realised by radically different perceptual contents.

Let's imagine, for instance, that the (perceivable part of) the electromagnetic spectrum was heard instead of seen. Different wavelengths could, in theory at least, give rise to differences in perceived pitch instead of colour. The distance of the objects on which the light reflects could, in turn, be perceived as a difference in volume. In principle, in other words, we could hear light. This, of course, is not saying that hearing light would be as useful as seeing it, nor that it is physically realisable by a hearing apparatus with anatomical resemblance to our own, but merely that it is logically possible. It is, as pointed out, possible in principle, which does not imply its practical feasibility. Similarly, we have to acknowledge the possibility that we could have developed radically different sense organs, sensitive to the same stimuli but mediating these stimuli in a different way, therefore yielding a different type of percept altogether.

It appears, indeed, that we have evolved particular phenomenal 'translations' of the external stimuli triggering our sensory apparatus. Our sensations are causally connected with external stimuli, but this connection is contingent. Different organisms could perceive the same stimuli in different ways.

2.4.3 The constructive role of the sensory apparatus

In order to extract useful data from the environment, our senses do not only mediate the stimuli to yield a particular percept, but also actively contribute to the content of the percept. Indeed, we do not merely receive images and sounds through our sensory receptors. Before it reaches our consciousness, it has already been processed by our mind in order to provide us with more information. Depth-perception appears in what was a two-dimensional projection on our retina, and sound is given a certain direction and distance. This information is based on automatic hypotheses that the mind constructs, using several cues. Without these mental constructions, our senses wouldn't be of much use. Vision, for instance, would be utterly useless if it merely provided us with a two-dimensional projection of colours. We need, on the contrary, to structure those rays of light in order to get a useful representation of the environment, enabling us to discriminate objects and judge distances. Auditory perception, in the same way, wouldn't be very helpful if we didn't (re)construct the direction of and the distance from the origin of the sound.

This constructive role of the senses, however, further divorces our perception from the external world causing it. Not only do we perceive but a (small) part of the world, and perceive this part in a contingent way, but, moreover, we add to the content of the percept. In other words, not all information we receive in perception comes from external stimuli. This paints a picture of perceptual abilities as detecting modules. Our senses detect particular, relevant elements in the environment, rather than duplicating the environment in our 'mind's eye'. Therefore, our perception of the world by no means mirrors the external world. We appear to be sensitive merely to a small range of elements in the environment (perceptual closure), to perceive these elements in a particular, contingent way, and to reconstruct the state of the world hypothetically, using several automatic cues.

This leaves us to conclude, that – just as any other organism – we are encapsulated in a particular 'experiential bubble'. Our perception of the world is both incomplete – only part of the potential external stimuli trigger our senses – and species-specific – the stimuli causing perception are perceived in a particular, contingent way and are completed with information that is not directly drawn from the external world.

3. Intuitive theories

3.1 Innate predisposition to individuate subject matters

Different subject matters require different kinds of explanations. When we see an animate creature moving fast we infer a motive behind this movement (e.g. it's fleeing from something or chasing something), but when we see an inanimate object moving, we interpret this movement in purely physical terms (e.g. the movement of a twirling feather is caused by the force that the wind exerts on it, or the rolling stone is moved by gravity). Similarly, when we encounter natural kinds (e.g. animal or vegetal organisms), we understand them – as Pinker (1997:314) puts it – 'in terms of their innards', whereas artefacts (e.g. chairs and tables) are understood in terms of the function they serve.

In this regard, Pinker (1997:315) argues that we are endowed with mental modules for dealing with objects and forces (intuitive physics), animate beings and other humans (intuitive psychology), artefacts (intuitive engineering) and natural kinds (intuitive biology), among other categories. Each category is approached in a different way. On a commonsensical level, it seems, indeed, absurd to try to explain the actions of other human beings in physical terms or, vice versa, to accord an intention behind the twirling movement of an air-borne feather.[1]

This predisposition to 'carve up the world' into different realms and to apply different intuitive theories to each of them appears to be at least partly innate. Infants distinguish between animates and inanimates (Spelke et al, 1995; Gelman et al, 1995) and treat them differently. They seem to grasp that animate creatures have an internal source of energy and intentions moving them, while inanimate objects can only be moved by external causes. Indeed, while they try to bring people to them by making noise, they bring objects to them by moving them physically (Pinker, 1997:322). Furthermore, this way of carving up the world and the different intuitive modes of thinking applied to them are similar across all human cultures (see Atran, 1998 on folk biology).

However, as Pinker (1997) points out, the fact that different ways of knowing are innate does not imply that knowledge is innate. It does not replace or minimise learning, but merely makes it possible. Indeed, Pinker continues, learning involves more than recording experience. We need to couch experiences so that they are generalisable in useful ways (315). In this sense, we must have a predisposition to interpret the behaviour of fellow human beings in terms of goals and values in order to make sense of it, or a predisposition to interpret our sensory data as made up by objects governed by physical forces.

Khalidi (2002), in this context, points out that innateness does not imply the presence of full-grown, innate ideas at birth, but requires environmental stimuli to develop it (252). Both 'nature' – i.e. innate predispositions – and 'nurture' – i.e. environmental stimuli – constitute our intuitive understanding of the world, which, therefore, come together when our cognitive

[1] This, however, is not saying that humans, in their creative frivolity, never think of inanimate objects as animated (cf. animistic religions) or of animate creatures as being reducible to inanimate particles (cf. physicalism), but that on a non-metaphysical level, non-critical level, people are predisposed to apply certain intuitions to certain categories.

architecture is fed input from the environment. While the relative importance of the nature versus nurture components of human behaviour is still heavily debated in the social sciences (cf. Pinker 2002), it is now commonly accepted that our mind must have at least a minimal role in structuring the perceptual input in order to yield knowledge. How, indeed, could one acquire knowledge based on observation alone? How could one make sense of any of it, without some form of (innate) cognitive mechanism that structures this perceptual input? Without the ability to, at least, detect similarities and differences among sensory data, there would be no way of classifying the input (e.g. of distinguishing night and day, animate and inanimate, animal and vegetal, etc.). These sensory impressions would be nothing but raw, unprocessed phenomenal experiences, remaining utterly unintelligible. Therefore, there is no escaping the conclusion that we need 'modes of knowing' – as Pinker puts it – to structure the input and generate knowledge. These innate mental predispositions, when informed by the perceptual input gathered from the environment, give rise to intuitive theories about the world. Those theories are uncritical, pre-scientific and shared by all human beings on a commonsensical level. They are often referred to as 'folk theories'.

Below, I will look at two deeply-rooted folk theories underlying our view of the world – folk physics and folk biology – and ask whether they correspond to their scientific counterparts.

3.2 Folk 'sciences'

3.2.1 Folk physics

Careful testing on infants has shown that they already show some basic appreciation of physical laws. Kellman and Spelke (1983), Spelke (1991) and Baillargeon (1991) have designed experiments on 3 to 8 month old children, to test their concept of objecthood and the laws that govern their interaction. In order to test this, they measure the looking time of the infant when confronted with either a possible or an impossible physical event (such as, for instance, an object passing through another or an object disappearing after being veiled). When infants consider something as an impossible physical event, their looking time will be considerably longer than when confronted with a possible event, which bores or 'habituates' them much faster, making them look away (Baillargeon et al, 1995:81).

Several conclusions emerge from these tests. First of all, infants possess the concept of objecthood. They appear to perceive an object whenever parts are moving together. Whatever constitutes an integrated whole, in this sense, is considered an object. (Kellman and Spelke 1983). They do not, in other words, perceive the world as ‘one great blooming buzzing confusion’, as James (1890:462) famously put it, but as an ordered whole made up of different objects. Indeed, when babies are shown two sticks poking up from behind a veil and moving in synchrony, they expect these to be part of a single object – i.e. to be attached. When shown that these are in fact two separate sticks they ‘express’ their wonder in their looking time. (Kellman and Spelke, 1983:493). However, when the sticks are not moving behind the veil, they expect them to be two separate objects (497-498).

Once established that infants perceive the world as made up of different objects, Baillargeon (1991) and Spelke (1991) went on to test their intuitions about objects and the physical laws governing them. They concluded that infants expect objects to be impenetrable by each other, to move along continuous trajectories and to be cohesive. Furthermore, they already ‘know’ that objects can only move each other by making contact. As Pinker (1997) points out, infants see objects, remember them and expect them to obey several physical laws. They have an understanding of a stable, lawful world, which they could never have acquired by simple induction (they are barely able to manipulate objects, they don’t see them very well, etc.) or through feedback from anyone else (they obviously can’t communicate). Therefore, they must be endowed with an innate predisposition to understand physical entities in a particular way (319).

As to the nature of this innate predisposition, however, opinions differ. Spelke (1995:45-51) argues that infants are endowed with core-beliefs or guiding principles when considering physical happenings. Baillargeon (1995:79-80), on the other hand, rejects the assumption that infants are born with substantive beliefs about objects, but claims that they are endowed with highly constrained mechanisms guiding their acquisition of knowledge of objects. In any case, whatever the underlying cause, both Spelke and Baillargeon established that intuitions about physical events have an innate basis – either in the form of core-beliefs or constrained knowledge gaining mechanisms – and underlie the way adults still intuitively think about these events.

This intuitive thinking, however, is not necessarily accurate. Indeed, our 'folk physics' is often at odds with scientific physics. As Pinker (1997) argues, our intuitions match Aristotelian physics – claiming that all objects strive towards rest – rather than Newtonian physics – claiming that all moving objects will continue indefinitely on their trajectory unless a force acts upon it. We spontaneously think that a moving object is impressed with an 'impetus' - a force acting upon it (e.g. the wind blowing on a leaf or a rock being thrown in the air) - until this impetus gradually dissipates and the object comes to rest. Similarly, when we say that 'the ridge keeps the pencil on the oblique writing table', we seem to imply that the pencil has a tendency to move (an impetus) and that the ridge overcomes this impetus by exerting a greater force, which not only gives us a distinctly unscientific account of events, but is also in direct contradiction with Newton's third law that states that action equals reaction (320). Furthermore, as Proffitt and Gilden (1989) have established, when it comes to more complicated motions, such as, for instance, wheels rolling down ramps, colliding balls or spinning tops, people's intuitions completely fail them in predicting the outcome.

However, as Pinker (1997) explains, the fact that the mind is non-Newtonian, is not surprising. In the real world, Newton's laws are masked by friction (from the air and contact with the ground). This friction slows everything down until it comes to a stop, making it very natural to conceive of objects as having an inherent tendency towards rest. Our intuitions have, indeed, not evolved to give us an accurate account of events, but merely to enable us to predict probable outcomes in our natural environment. The same reason explains our failing intuition when it comes to complicated motions, as these complicated motions are very unlikely to happen in natural environments (321).

3.2.2 Folk biology

People everywhere have deep-rooted intuitions about natural kinds, such as animals, plants and minerals. According to Atran (1998), we are endowed with a predisposition to think about fauna and flora in a highly structured way. Indeed, we divide the natural world into a complex taxonomy which incorporates different groups, each further defined in different levels of subgroups (e.g. a lion is an animal, a mammal and a cat). Atran argues that these taxonomies are widely shared across all cultures and eras, and are therefore much less arbitrary than the assembly of, for instance, entities in cosmology, artefacts or social groups (547). As Simpson (1961:57) puts it, classifying animals and plants into basic groupings is ‘quite as obvious to [the] modern scientist as to the Guarani Indian’.

This innate predisposition manifests itself – as can be expected – at an early age, where children not only distinguish between animals and non-animals (Atran 1998:549) – corresponding to the animate–inanimate divide pointed out above - but also between plants – inanimate but organic – and non-living things (Gelman and Wellman, 1991). With respect to the natural, organic world, as mentioned above, human minds intuitively distinguish different realms, each realm itself divided up into groups and sub-groups.

This predisposition to classify the organic world according to a complex taxonomy arises from an intuition of a hidden trait or essence that members of the same group share with each other (Pinker, 1997:323). As Atran (1995:219-220) explains, the nature of this essence is initially unknown to children, but is already presumed. We are, in other words, predisposed to look for essences when dealing with the natural world. Interestingly enough, these essences are not based on visible similarity but on what is presumed to be the underlying constitution (Pinker, 1997:324). It rests on the assumption that natural kinds have an underlying causal nature, uniquely responsible for their typical appearance, behaviour and ecological niche. (Atran, 1998:548). It is also what makes the caterpillar the same animal as the butterfly it develops into. Although its appearance might be radically changed, we are still inclined to reason that its underlying constitution remains constant and causes the organism to develop into a new form. As Atran (1998:548) puts it, the essence maintains the organism’s integrity even as it causes the organism to grow, change and reproduce, in which process it passes down its essence to a new organism.

Gelman and Wellman (1991) argue that this essentialism develops in children at a certain age. For instance, when asked what happens when doctors take a tiger, bleach its fur and sew on a mane (making it look like a lion), five year olds typically say that it is now is a lion, while seven year olds say it’s still a tiger, pointing to the fact that they attribute an identity to an animal based on its innards or its hidden essence, rather than its external appearance. Pinker (1997:327) argues that this essentialism cannot be learnt. Indeed, children have not taken biology at that age and parents commonly do not describe animals in terms of their innards or internal constitution. Furthermore, children never develop essentialist thinking about artefacts. (Pinker, 1997:326). A table that is taken to pieces and reassembled as a chair is for all people at all ages now a chair. In this regard, it seems established – as pointed out above – that people treat natural kinds and artificial kinds differently, intuitively according essences to the former, while describing the latter in terms of the function they serve.

Once again, as is the case for folk-physics, our intuitive grasp of the natural world is contradicted by modern science. Indeed, while this essentialist thinking was still the basis of pre-Darwinian scientific taxonomy, known as the 'Linnaean classification system', Darwin's evolutionary theory proved this intuitive thinking wrong. Species do not possess immutable essences, but are themselves subject to change. Birds evolved out of reptiles and elephants out of rodents; all 'essences' appear to be nothing but temporary forms of species in adaptation to specific environmental conditions. Species do not fit into classes and subclasses because of their essence, but because of the relative proximity of their common ancestors.

In a broader perspective, this categorical thinking is criticised by Lakoff (1987), who claims that there are no clear-cut categories applicable to the world at all. A prima facie unambiguous category, such as, for instance, mother – which could be defined as female genitor or parent – instantly becomes more problematic when borderline cases have to be accommodated. What about an adoptive mother, the woman ceding a donor egg, or the case in which a 'surrogate mother' has a fertilised egg implanted in her uterus and gives birth? If you say the origin of the egg counts, women having donor eggs implanted are not mothers. If, on the contrary, you say that giving birth makes one a mother, what about 'surrogate mothers' who 'rent' out their uterus to bring somebody else's child into the world? If you want to keep parental care as the only condition, you include adoptive mothers, but what about the biological mother? In whatever way you twist and turn the issue, there seems to be no necessary and sufficient condition to define the category. Similarly, biological categories notoriously resist an unambiguous definition. The category of 'fish', for instance, appears impossible to define[2], and the category of mammals became problematic when naturalists

[2] Fish, as Pinker (1997:311) points out, 'do not occupy one branch in the tree of life. One of their kind, a lungfish, begot the amphibians, whose descendants embrace the reptiles, whose descendants embrace the birds and the mammals. There is no definition that picks out all and only the fish, no branch of the tree of life that

stumbled upon the Platypus, an egg-laying, breast-feeding species living in remote corners of Australia and Tasmania. This, of course, is the result of evolution, where species drift in and out of larger categories, new categories are formed, and old ones disappear.

However, while our intuitive way of structuring animal and vegetal organisms in fixed categories might be scientifically inaccurate, it is nevertheless very useful in dealing with the environment. Our life spans are, indeed, considerably too short to have to take the evolution of species into account. Considering organisms as endowed with immutable essences, therefore, works perfectly well. Furthermore, intuitive theories enable us to assign properties to unknown but 'similar' elements in the environment. A category of poisonous animals, for instance, might not be clearly defined and might even mistakenly include species that are not poisonous but that have a lot of characteristics in common with species that are. This, however, does not change the undeniable advantage of having an approximate (and, in this case, conservative) classification. Indeed, as Pinker (1997:312) rightly points out: 'systems of rules are idealizations that abstract away from complicating aspects of reality'. They evolved to provide us with a workable framework of the world around us, not to accommodate all ambiguities and complexities.

3.3 Conclusion

3.3.1 Innateness

Our intuitive grasp of the world appears to be founded on innate predispositions. Both tests on infants enquiring about their uninformed expectations about states of the world (cf. Spelke, 1995 and Baillargeon, 1995) and comparative anthropological research (cf. Atran, 1998), show that the human mind is endowed with an innate predisposition to make sense of its environment. This predisposition underlies the way we carve up the world, explain events and predict probable outcomes. While this view has traditionally been opposed by empiricist theories, which argue that the mind is nothing but an empty shell, it is now commonly accepted that the mind is endowed with genetically based modes of knowing. As pointed out

includes salmon and lungfish but excludes lizards and cows.'

above, tests on infants and comparative anthropological research point conclusively towards a shared, innate, intuitive view of the world.

3.3.2 Scientific inaccuracy

This intuitive view, however, is not an accurate and objective one. Indeed, both the way we split the world into different categories and the different intuitive theories we apply to those categories appear at odds with a scientific account of the world. Regarding our predisposition to individuate subject matters, it appears that even animate and inanimate or organic (living) and inorganic (lifeless) entities – giving rise to a very different intuitive conception – are not accounted for in radically different ways by modern science. Indeed, science tells us that all matter – which includes, of course, animal and human organisms as well as rocks – is made up of molecules of atoms. Even animate organisms, in this sense, obey purely physical laws, where its movements are caused by nothing other than atomic interaction (within the brain, from brain to muscles, and so on). This physicalistic conception of living organisms clashes with our deeply-rooted intuition that living organisms possess a quality that is absent in other matter. It is, therefore, not surprising that the premise that everything is physical has notoriously awakened vivid criticism throughout history and is still heavily debated today with regards to the mind in general and consciousness in particular.

Furthermore, as pointed out above, our conceptual framework providing us with an intuitive grasp of physical happenings – ascribing an 'impetus' to moving objects, assuming that every object's natural state is rest, and so on – contradicts Newtonian physics, let alone Einstein's account of the universe. Similarly, our essentialist categorising of the organic world is in direct contradiction with the theory of evolution. Species are not endowed with immutable, internal 'identities', but change over time. Organisms, in this sense, are only part of species and larger groups because of genetic similarities due to the proximity of a common ancestor, not because of a shared essence that resists time.

3.3.3 Species-specificity

While our intuitive grasp of the world obviously fails in terms of scientific accuracy, it nevertheless provides us with a very useful framework, specifically tailored to our needs. We do not need to grasp how celestial bodies interact with each other or how matter is constituted by atoms to get by on this planet. Folk physics, with its naïve, uncritical conceptual basis for explaining and predicting events in this world, provides us, nevertheless, with all the information we need to function. Indeed, both the macroscopic and microscopic dimension do not concern our ability to manage our surroundings. We are confined to the 'medium dimension' that we occupy – somewhere in between electrons and galaxies – on the particular planet we live, with its gravity, air resistance, and so on. Those are the only elements that our intuitive grasp takes into account, yielding an erroneous, biologically biased view, shaped to fit our 'ecological niche'.

Similarly, in the biological realm, we do not have to be evolutionary geneticists to predict the behavioural traits of animals, the medicinal properties of certain plants, or the environment in which they will thrive. In order to do this, our erroneous essentialist conception, carving up the natural world into a complex taxonomy, is more than adequate. It provides us with a very useful framework to deal with our environment, enabling us to predict properties of an unknown element based on its similarity with known elements. Without this ability to categorise, our knowledge could not extend beyond the realm of elements we are actually acquainted with and would leave us utterly unprepared for any new element.

Therefore, our intuitive theories about the world are thoroughly species-specific. They are tailored to the needs of our particular species in our particular environment. They yield a biologically biased view of the world that – while obviously mistaken from a scientific point of view – enables us to function in the world. This, of course, is the result of an evolutionary process adapting organisms to meet their biological needs in their environment.

4. Why do we perceive and make sense of the world the way we do?

Both our perceptual and cognitive abilities – the mental mechanisms underlying our intuitive assumptions about the world – are the outcome of an evolutionary process, driven by the blind mechanism of natural selection. The theory of evolution by means of natural selection, as formulated by Charles Darwin (1859) and Alfred Russel Wallace (1858)[3], states that all species evolve through the selective retention of random mutations and genetic recombinations occurring in their gene pool.[4] Variations that increase an organism's chance for survival and reproduction in their particular environment will soon spread over the entire population, because they are passed on by organisms with a higher than average chance of reproducing. Variations that negatively affect chances for reproduction will, by that same rationale, soon be weeded out of the gene pool. Therefore, species evolve and can evolve more complex abilities. This is how a single cell, constantly copying itself, eventually evolved – through copying errors that proved beneficial – into the mind boggling variety of flora and fauna flourishing on this planet. This process is called natural selection.

Selection is the product of what Darwin calls 'Malthusian pressures'[5], meaning that more individuals are born than can survive and successfully reproduce. This results in a struggle for life. Organisms with beneficial characteristics will get the upper hand in this struggle and their genes will spread throughout a species' gene pool. With regard to this struggle for life, Darwin rhetorically asks when introducing the concept of natural selection:

---

[3] Wallace sent Darwin a letter, with a fairly brief account of the theory of evolution by means of natural selection while Darwin was still postponing publication of the Origin of species, after decades of research.

[4] More recently, this Darwinian account of evolution by purely environmental (external) selection has been criticised by Riedl (1984a) and Wuketits (1990), among others. The new view defends a 'systems theory of evolution', stating that organisms are not simply moulded by their environment, but are active systems developing better living conditions. Organisms, in this regard, evolve both by external and internal (intraorganismic) selection. Evolution, in other words, is constrained and influenced by the structures and functions of organisms themselves. There is a feedback flow of information from phenotype to genotype. It is argued, that without this feedback and internal regulatory constraints informing the genotype about phenotypic requirements, random mutations could not account for the order of systems and subsystems in the organism. All forms of life would break down in chaos. (Riedl 1984b, Wuketits 1990). However, the fact that there are constraints from within the organism, along with external conditions, does not change our premise that evolution is a blind process, gradually adapting organisms to particular needs in a particular environment based on selective retention of random (undirected) variations.

[5] The English economist, Thomas Malthus, stated that available living space and food supply could never keep up with the increase of population. Populations have a tendency to grow much too fast for the increase in resources, which inevitably leads to a struggle for existence, leading to war, epidemics and famine. (Malthus 1798).

> [C]an we doubt (remembering that many more individuals are born than can possibly survive) that individuals having any advantage, however slight, over others, would have the best chance of surviving and of procreating their kind? On the other hand, we may feel sure that any variation in the least degree injurious would be rigidly destroyed. This preservation of favourable variations and the rejection of injurious variations, I call Natural Selection (Darwin, 1859:80-81).

The origin of our perceptual and cognitive abilities resides in this selective process, slowly shaping organisms by retaining favourable variations and rejecting injurious ones. They evolved to increase an organism's chances for survival and reproduction. In this sense, every organism's mental abilities gradually evolved together with their sensory abilities, pushed by the selective advantage of enabling the organism to extract more useful elements out of its environments and avoid more harmful ones. This implies two things: firstly, that our perceptual and cognitive abilities are tailored to the *particular needs* of our ancestors, and secondly, that these abilities have only been selected for when they enabled our ancestors to *behave* in ways that increased their chances for survival and reproduction. Abilities that do not translate into fitness enhancing behaviour will not be selected for.

4.1 Perception

Perception might not yield an objective and complete representation of reality, but it nevertheless seems inconceivable that evolution provided us with a simply useless realm of awareness. As explained above, the mechanism of natural selection entails that it evolved to increase our chances of survival and reproduction. In this regard, our perceptual faculties must have evolved to provide us with information eliciting 'appropriate' behaviour in correspondence with our surroundings. Vision allows us to avoid obstacles, to stalk prey, and to find our way home. Auditory perception, on the other hand, warns us of approaching danger and allows us to communicate at a distance. The chemical senses, finally, cause us to avoid eating putrid substances, while the sensation of pain stops us from damaging our body, to name but a few examples.

Indeed, gradual adaptations in our evolutionary history have provided us with an extremely useful realm of information which we draw from our environment, guiding our actions in appropriate ways. This leaves us with a very useful ability that is, on the one hand, irremediably particular, on the other hand, selective. Vast realms of potential stimuli, in this context, remain beyond our awareness, simply because their detection has not been selected for. Boynton (1979) points out what he considers an apparent paradox:

> Considered within the framework provided by the theory of evolution, it is perhaps not particularly difficult to understand why we should be so limited and yet at the same time so gifted. Our perceptual abilities have presumably evolved, along with our other capacities, according to what is important in the external environment from the standpoint of adaptation and survival. If, say, static fields of magnetic flux carried significant information for us, we probably would have evolved magnetoreceptors (Boynton, 1979:44).

Every organism's perceptual abilities, in this sense, detect the particular elements in the environment that provide an input which increases an organism's chances for successful reproduction. These elements are different for every species. Therefore, different species perceive the world differently. Their senses are triggered by different stimuli and they often perceive similar stimuli in different ways.

4.1.1 Detecting 'affordances'

Gibson's (1979) notion of 'affordances' is illuminating in this context. An affordance of the environment is an element or aspect of it that is relevant for an organism. It is something the environment provides the organism with, either in a beneficial way or a harmful way. The same environment, therefore, can hold different affordances for different species (127). An impala for instance, will have as affordances the grassy plains, the watering holes, and the big carnivorous cats in its environment. A fly, living in the same environment, will, of course, have very different affordances, such as the presence of dung on which it feeds, birds and reptiles that feed on it, and so on. On an even more basic level, an affordance of the impala would be the terrestrial surface on which it moves while that of the fly would be the air through which it moves.

Indeed, every species has evolved to meet nature in a certain way. Each of them has carved its ecological niche, developing a certain way of life using particular affordances in its environment, and allowing the species to subsist. An ecological niche, in this sense, is a set of affordances. Species develop adaptations to these affordances. Their biological constitution is shaped by natural selection to maximise the benefits it can reap from its environment (in order to keep the organism safe, to nourish it and to reproduce), and to minimise the chances of encounters with elements of the environment that threaten its survival and/or chances of reproduction (e.g. preying animals, poisonous substances, competition for reproduction, etc.) (129).

These adaptations, however, do not only include physiological adaptations (such as the extremely light weight that is divided over long legs which enables the water bug to stride on water, the hydrodynamic shape of fish or the camouflaged colours of so many animals). Perceptual abilities have also evolved in relation to those affordances. Each species is endowed with a particular way of detecting its affordances and with a way of translating this detection into action. The success in benefiting from positive affordances and avoiding negative ones is the only reality in the light of which natural selection has shaped the perceptual abilities of every species. In this context, our perceptual experiences are shaped to provide us with a useful picture of the particular environment in which we have spent most of our recent evolutionary history.

4.1.2 Using the perceptual input

Detecting affordances, however, as mentioned previously, is not enough for natural selection to shape the perceptual apparatus. This detection needs to elicit fitness enhancing behaviour. A species only develops the ability to perceive affordances insofar as it possesses the ability to 'translate' this input into 'appropriate' behaviour. As Matthen (2005) points out, there is an interdependence of perception and use that explains the specialisation of perceptual abilities of organisms. Indeed, if an organism does not translate the perception of particular features of its environment into action, it will not evolve the ability to perceive it. Since perceiving these features would not enhance its chances for reproduction, the genetic variations allowing for this ability will not be selected for. Reciprocally, an organism can only develop behavioural strategies if it has evolved the necessary perceptual abilities to gather the required information for these strategies (223). Therefore, both perceptual abilities and either cognitive abilities (in the case of humans and higher scale animals) or effector cues[6] (in the case of lower scale animals), co-evolved to 'match' an organism's bodily actions with the relevant changes in its environment.

Perception, in this light, is used differently by different organisms. It has evolved, as David Marr explains, to serve different purposes.

> A pigeon uses vision to help it navigate, fly, and seek out food. Many types of jumping spider use vision to tell the difference between a potential meal and a potential mate. (…) The frog (…) detects bugs with its retina, and the rabbit retina is full of special gadgets, including what is apparently a hawk detector, since it responds well to the pattern made by a preying hawk hovering overhead. Human vision, on the other hand, seems to be very much more general, although it clearly contains a variety of special-purpose mechanisms that can, for example, direct the eye toward an unexpected movement in the visual field or cause one to blink or otherwise avoid something that approaches one's head too quickly. (...) The general point here is that because vision is used by different animals for such a wide variety of purposes, it is inconceivable that all seeing animals use the same representations; each can confidently be expected to use one or more representations that are nicely tailored to the owner's purposes (Marr, 1982:32).

In this regard, since perceptual abilities are used for different purposes, and since different elements in the environment matter for satisfying those purposes, the perceptual input that organisms gather from their environment can differ radically. Human perceptual abilities are no exception, as they are the outcome of our evolutionary history.

[6] Effector cues are immediate behavioural reactions triggered by a particular perceptual sensation. A good example is the reflex of a tick to drop from its leaf when its perceptual faculties are triggered (by sensing butyric acid and the exact temperature of 37° Celsius, indicating the presence of a mammal).

4.1.3 Human perception

As Matthen (2005:181) points out, homo sapiens inherited from its primate ancestors a special kind of colour vision that facilitates the detection of fruit among foliage (thanks to a particularly sharp contrast between red, orange and green). Similarly, our stereoscopic vision, enabling depth-vision – crucial for swinging from one branch to the other – was selected for in our tree-dwelling ancestors. However, our hominid ancestors, while still closely linked to their primate ancestors, swapped an existence as arboreal frugivores for that of hunters. Therefore, according to Campbell (1985:370), our perceptual abilities must have evolved. Indeed, hunting animals' perceptual abilities are specialised to identify mobile objects (i.e. their prey), while frugivores' perception principally serves to identify static objects (plants containing fruit). Kortlandt (1965), in this context, recorded that when zoo chimpanzees receive food at a location at which they are not accustomed to receive it, they react as though it is something strange and often refuse to eat it. This is, of course, not the case for carnivores. A cat will find and eat its food wherever it is placed.

However, the main difference, as Campbell (1985) notes, between homo sapiens and other hunting animals, is that the latter mainly hunt by olfactory perception, while we – having inherited a highly developed sense of vision from our primate ancestors – 'rebooted' our primary sense – i.e. vision – to serve the purpose of hunting. We evolved the ability to identify objects on the move, visually, without any reference to a fixed position in the environment, and became the first land animals to hunt almost exclusively by sight (370). The result is a particular visual representation of the world, different from any other organism, as the outcome of both our primate heritage and its adaptation to new selective pressures associated with hunting.

Similarly, our auditory faculty has evolved to discriminate and analyse sound to a much greater extent than that of any other species (Campbell, 1985). This ability, allowing for human spoken language, is the result of the evolution of the auditory apparatus in primates, where hearing was mainly concerned with detecting elements from the social environment – i.e. from other members of the species – rather than the general environment. This, once again, yields a perception of the world that is very different from that of most other species, who use – for the most part at least – hearing as a general detection mechanism in the environment, often hearing fainter sounds, but not discriminating to the same extent between different kinds of sound (247-248).

4.1.4 Umwelt

Every species is, therefore, encapsulated in a particular realm of awareness as the outcome of its evolutionary history. Jacob von Uexküll (1909) coined the term 'umwelt' to describe this perceptual encapsulation. All species, he states, have their own umwelt - their own ambient reality - that is both the product of their environment and the way of perceiving this environment. In this sense, every living being – including ourselves – is 'locked up' in a particular phenomenal reality. The world, Uexküll rightly states, is never perceived objectively. Its perception is always mediated by the particular functioning of the sense organs that an organism has acquired in its evolutionary history.

4.2 Cognition

4.2.1 Increasing control over the environment

While our perceptual abilities have evolved to provide us with useful information about the environment, our cognitive abilities, underlying our mental predisposition to structure and understand the world in a certain way, have evolved to interpret this input in ways leading to 'appropriate' behavioural responses. As perception without cognition remains utterly unintelligible, both abilities have necessarily co-evolved in organisms to match their actions to the constantly changing states of the environment. Just as our perceptual abilities have only evolved to yield an input that is useful for the needs of our particular organisms, so have our cognitive abilities only evolved to provide us with a grasp of our environment that increases our chances for successful reproduction. They have evolved in order to increase our control over the environment, enabling us to extract more useful elements from it and to avoid more harmful ones.

The evolution of mind, in other words, is directed by the selective advantage of the increasing ability of organisms to use the environment to their benefit. In this context, Dennett (1995) proposes an account of different levels of mind, each level providing the organism with more efficient ways to cope with the environment. As each new level is reached, it makes it possible for the evolving organism to 'construct' the next one, enabling it to interact with its environment with increasing efficiency. Dennett calls this the 'tower of generate and test' (373).

At the basic level of Dennett's tower you have 'Darwinian creatures', where each species has evolved through natural selection in adaptation to its environment, but is entirely 'hardwired'. The organisms cannot adjust to their environment - only the species themselves can adjust over the generations by retaining advantageous variations in their gene-pool (374).

At the second level you have 'Skinnerian creatures', named after the behavioural psychologist B.F. Skinner, who pointed out the possibility of acquired behaviour (not genetically fixed) through the process of conditioning. These creatures can try out different actions in response to their environment and retain, by means of trial and error, the actions that worked. This, however, is still a pretty laborious and perilous process. Indeed, trying out random options can take a while to yield a fitting one and organisms can pay with their life for trying out risky errors (374).

Therefore, the next level involves a mental selection of those endless possibilities, avoiding the risky activity of testing out too many errors. Organisms capable of this feat are called 'Popperian creatures', since, according to Dennett, as 'Sir Karl Popper once elegantly put it, this design enhancement permits our hypotheses to die in our stead' (375). This level requires mental representations of actions in a particular situation and their probable outcome. This is, according to Dennett, not an exclusively human ability, and not even an exclusively primate or even mammalian ability. Birds, reptiles and fish are all endowed with this capacity to use information from their environment and mentally select their behavioural options before putting them into action (376).

The last level, on the other hand, is attained to only a very small extent by other animals. It is referred to by Dennett as the level of 'Gregorian creatures', after the British psychologist Richard Gregory, who emphasises the role of information, transmitted non-genetically by members of a group in what Dennett calls, 'the creation of smart moves' (377). We humans are endowed with language, an abstract framework of symbols picturing reality. Therefore, we can profit from detailed information thought up by our fellow human beings. The invention of writing, furthermore, exponentially increased the available amount of information. Writing can be stored on papyrus, clay tablets, in books and now even on hard drives, and this eliminates the necessity of mouth to ear transfer, making it accessible to a much larger number of individuals and to future generations.

The picture Dennett paints of the evolution of mind is one pressured by the biological advantage of gaining increasing control over the environment. It shows that the human mind – even in its vertiginous complexity – can be accounted for in naturalistic, evolutionary terms, by the hand of natural selection. There is, indeed, an obvious advantage to be gained at each new level of mind, all the way up to the Gregorian mind, allowing for human culture.[7]

4.2.2 Adaptive problems

While Dennett shows that the evolution of mind in general is pushed by the adaptive advantage of providing the organism with an increasing ability to use the environment to its benefit, he does not provide an account of the particular elements in the environment that our minds have evolved to control. This is the field of evolutionary psychologists. Tooby and Cosmides (1992, 2003), along with others (e.g. Pinker, 1997), fill this void. They argue that our mental abilities have evolved to meet certain challenges in our environment - so-called adaptive problems. Those adaptive problems, best described as reproductive opportunities or

[7] Miller (2000), however, argues that a purely survival oriented theory of selection can never account for some abilities of the human mind. He points towards a different, but equally important driving force of evolution: sexual selection. Arguing that just as the tail of the peacock did not evolve to provide the male birds with increased survival chances – quite to the contrary, the tail makes them more exposed to predators – some mental abilities in humans, such as, for instance, art, music, drama and comedy, cannot be accounted for by survival enhancing selection. These abilities, according to Miller, are the product of sexual selection through mate choice. They developed because pre-hominids endowed with them were preferred mates and therefore were able to pass down their genes to a greater extent. In this regard, these mental capacities are best seen as a 'sexual ornament', in the same way as the peacock's tail or the feathers of the paradise bird.

obstacles created by enduring conditions in the environment, include a multitude of domains, such as winning social support from group members, detecting cheaters, remembering the location of edible plants, hitting game animals with projectiles, identifying objects, assessing the character of self and others, acquiring language, thwarting enemies, and so on (Cosmides and Tooby, 2003:59).

According to Tooby and Cosmides, adaptive problems have two defining characteristics. Firstly, they are recurrent conditions in the environment, meaning that most individual ancestors encountered them, and that these conditions reappeared again and again over the evolutionary history of our species. Secondly, they are 'a subset of enduring relationships' that could, in principle, be exploited by a property of an organism in order to increase its chances for successful reproduction. Those problems could, in other words, be met with an adaptive response, thereby endowing the organism with an advantage in terms of reproductive success. (Cosmides and Tooby, 2003:59).

Adaptive problems are, by their very nature, species-specific. Since different species have different sets of affordances (cf. Gibson, 1979), the reproductive opportunities and obstacles they encounter will, of course, be very different. Therefore, just as perceptual abilities do not evolve in an organism in order to provide it with a neutral window on the world, but rather in order to detect what matters for that particular organism, so cognitive abilities evolved in the light of a particular set of challenges that the environment holds for a certain type of organism. This results in a particular species-specific way of framing the perceptual input – a particular grasp of the world – directly related to the adaptive problems encountered throughout the evolutionary history of a species.

4.2.3 Particular intuitive assumptions about the world

Our innate, cognitive predisposition to carve up the world in a certain way and our various intuitive theories about the world are, therefore, shaped in the light of the particular adaptive problems faced by our ancestors. They provide us with a framework that enabled our ancestors to thrive in their environment, increasing the efficiency to use it to their benefit. They do not, however, as indicated in the previous section, provide us with scientifically correct assumptions about the world. Indeed, they are not shaped to enlighten us with the true, objective structures of reality, but to endow us with assumptions that work to our benefit in terms of survival and reproduction in the particular environment we have occupied throughout our evolutionary history. Our intuitive framing of the perceptual input gathered through our senses is, therefore, both fallible and thoroughly species-specific.

5. Conclusion

This chapter set out to investigate the way in which we intuitively view (i.e. perceive and understand) the world. In order to do this, I looked at perception and cognition. Adopting a naturalistic perspective, I reasoned from the prevailing scientific account of those abilities, casting them in an evolutionary context.

Regarding perception, I argued that we perceive but part of the external world (perceptual closure), that we perceive this part in a particular way (the contingent relation between perceptual cause and perceptual content) and that our sensory apparatuses actively contribute to the content of the percept. Therefore, it follows that what we perceive is a radically incomplete and species-specific representation of the world.

Regarding cognition, I argued that we are endowed with an innate predisposition to interpret our perceptual input in a particular way. Our mind intuitively carves up the world and applies different modes of thinking to different categories. This results in a variety of intuitive theories about the world, which are incorrect from a scientific point of view.

Both our perceptual and cognitive abilities are, indeed, the product of evolution by natural selection, entailing that they have evolved in response to particular aspects of the environment which mattered to our particular organisms in terms of survival and reproduction throughout our evolutionary history. These abilities evolved to provide us with the right input and the right way of processing this input for appropriate action schemes in our natural environment.

This leaves us with both a particular, incomplete input (perception) and particular intuitive assumptions about the world (cognition).

Therefore, our intuitive view of the world, as the result of both the senses and the cognitive predispositions with which we are endowed, is a thoroughly particular one. Indeed, we are perceptually closed to parts of the world, and what we perceive comes to us in a particular way. Furthermore, the way we interpret this perception is the result of innate predispositions providing us with intuitive theories about the world that are at odds with modern science. We are, in other words, encapsulated – on a commonsensical level – in a contingent, species-specific scope on the world, as a direct consequence of the evolutionary road we took. This view is not imposed upon us merely by the structures of the world (as it is based on our particular and incomplete perception of the world and our intuitive assumptions, which are contradicted by modern science) and is not the only possible representation (meaning that organisms with a different perceptual and cognitive make-up could view the same world differently).[8] Indeed, in the same way as every other organism on this planet, we grasp the world - on an intuitive level – in a specific way that does not yield an accurate, objective and complete representation of it.

[8] Cf. Definition of 'particular' in the introduction of this chapter.

Chapter 2: The Lorenzian fallacy: deducing truth from functionality

1. Introduction

In this chapter I will outline the argument of Konrad Lorenz (1941, 1973) which states that, since our perceptual and mental abilities are the product of natural selection, the representation of the world they provide us with must be accurate. This argument, very influential in evolutionary epistemology[9], contradicts the conclusion of the previous chapter that our intuitive representation of the world is particular and therefore does not correspond in any straightforward way to the structures of the external world.

I will argue that this central argument of evolutionary epistemology is fallacious. In order to support this position, I will look at Alvin Plantinga's (1993) argument against evolutionary naturalism, pointing out the fact that isolated beliefs remain outside of natural selection's radar. Furthermore, I will ask whether rationality is automatically fitness enhancing. Following Stich (1990), I will argue that in some cases natural selection will not lead to rationality. In this context, I will point at cognitive biases that we appear to succumb to, and throw an evolutionary light on them, as well as at non-cognitive purposes that belief-forming serves.

The mechanism of evolution by natural selection will also be taken into account. Drawing from the fact that variation is generated randomly and the selection of particular traits and abilities are both dependent on an organism's evolutionary history (path-dependency) and its particular environment, I will argue that this does not guarantee us the accuracy of our gene-based representation of the world. Moreover, I will point towards the existence of genetic drift and genetic 'hitch-hiking' which allows for certain traits and abilities to spread, regardless of their adaptive value.

[9] Evolutionary epistemology, in the context of this dissertation, refers to the research program, labelled 'evolution of epistemic mechanisms' by Bradie (1986) (cf. Introduction).

Finally, I will argue that the Lorenzian argument – inferring the accuracy of our mental representation a priori from the fact that there must be a connection between adaptations and the environment, and that our mental faculties must be adapted – fails in principle because of a lack of external perspective.

2. Lorenz's evolutionary epistemology

At the dawn of World War II, Konrad Lorenz (1941), an Austrian biologist, published an article that proposed a radical new perspective on traditional epistemological questions by situating them in a Darwinian context. In this article, Lorenz attempts to provide an evolutionary account of the working of our cognitive mechanisms and points out the consequences that this holds for our epistemic endeavours. While similar approaches have been hinted at by Herbert Spencer (1855), Charles Darwin (1871) himself and famous pragmatists such as William James (1890) and John Dewey (1910), Lorenz was the first to elaborate this view and to offer empirical evidence in support of it.

Despite its importance and implications, however, it remained largely unknown for over three decades. The world was at war, Lorenz was Austrian, and to make matters worse he joined the Nazi party and published pseudo-scientific writings supporting Nazi ideology. However, more than three decades later, Lorenz's groundbreaking work in the field of biology was recognised, and he was awarded a Nobel Prize, together with his colleagues Nikolaas Tinbergen and Karl von Frish. He officially apologised for his previous political opinions and was embraced by the world as one of the founders of a new science: 'ethology' - the study of animal cognition and behaviour.

In this section, I will outline what I refer to as 'the Lorenzian argument', at the core of evolutionary epistemology[10], stating that:

[10] When referring to evolutionary epistemology, I aim at those theories reasoning about human knowledge from the premise that the cognitive abilities underlying it are the outcome of biological evolution. As pointed out (cf. General introduction), a distinct program also runs by the name of evolutionary epistemology, drawing an analogy between the evolution of ideas or theories and biological evolution by means of blind variation and selective retention. These theories, developed by thinkers as Popper (1972) and Campbell (1974), are not our concern here.

> Since our perceptual and cognitive apparatus, providing us with a representation of the world, are the result of evolution by natural selection, the structures of this representation must – at least approximately – match the structures of the world itself, because being endowed with an accurate representation of the world increases one's biological fitness and therefore the perceptual and cognitive mechanisms allowing for this accurate representation must have been selected.

2.1 Ontogenetic a priori as a phylogenetic a posteriori

2.1.1 Adaptations as representations

Lorenz (1941) starts his argument by stating that all features of organisms 'mirror' nature. Indeed, he claims, the process of natural selection has adapted these features to the environment in which the organism lives. The eye reflects the laws of optics in its very structure, the shape of a fish reflects the laws of hydrodynamics, and the anatomy of a bird's wings those of aerodynamics. All characteristics of an organism, shaped by the process of natural selection, reflect the elements of the environment to which they are adapted.

In this regard, Peter Munz (1993) views organisms as 'embodied theories'. The organism, he claims, incarnates theories about the environment - it reflects 'knowledge' about the environment. According to Munz, however, this does not mean that the organism mirrors or even describes those parts of the world to which it is adapted, but merely, that some properties of the organism disclose information about certain properties of the environment. The hydrodynamic shape of fish does not describe water, let alone mirror it, but tells us something about the properties of water in relation to locomotion. This representation, of course, is neither verbal nor conscious. It is expressed through an organism's anatomical structure and behavioural programming.

In this sense, it is embodied, not disembodied – meaning that it cannot be separated from the organism expressing it. When the kingfisher dives into the water to catch a fish, it does so by correcting the angle of refraction. However, it does not *know* Snell's law – calculating the angle of the refraction of light when it passes from air to water – but, that knowledge is, unconsciously and in a non-verbal, non-mathematical way, present in its genetic constitution, triggering a particular action-scheme when it spots a fish from above (154-158). This is what is meant by 'incarnated knowledge' or 'embodied theory'. Adaptations, in this perspective, are representations of aspects of the world.

2.1.2 A non-arbitrary a priori

Lorenz (1941) argues that human perceptual and cognitive structures are no exception. Our mental a priori, structuring data into (conscious) knowledge, has evolved through natural selection in the same way as every other characteristic or ability of any organism. This implies, according to Lorenz, that they are not arbitrary. Just as any other feature of an organism reflects or represents the structures of the elements of the environment to which they are adapted, so do our perceptual and cognitive abilities.

> Just as the hoof of the horses is adapted to the ground of the steppe which it copes with, so our central nervous apparatus for organizing our image of the world is adapted to the real world with which man has to cope. Just like any organ, this apparatus has attained its expedient species-preserving form through this coping of real with the real during a species history many eons long (Lorenz, 1941:124).

In other words, our perceptual and cognitive systems are shaped by natural selection to cope with the external world. Therefore, he argues, it must provide us with an accurate representation of reality, allowing us to make true inferences about what really is the case. If, indeed, our mental abilities did not represent reality in a truthful way, Lorenz argues, they could never have evolved. Organisms endowed with perceptual and cognitive structures that did not lead to a truthful representation of reality would have less than average chances for survival and reproduction. Therefore the genes responsible for these inaccurate representations of reality would soon be weeded out of the gene-pool. Natural selection shapes every organism to 'fit' their environment - similarly it has shaped our mental a priori to 'fit' the world around us, i.e. to yield a truthful representation of objective reality.

In this context, Lorenz casts the Kantian a priori – our innate perceptual and analytical structuring of external data – in an evolutionary light. Lorenz agrees with Kant that we are endowed with a mental a priori - we do not merely absorb reality as it is, but actively structure it, both through perception and understanding. He denies, however, that these perceptual and analytical structures are necessary and universal. According to Lorenz, they are the product of evolution by natural selection and are therefore 'tuned' to the external constraints of the world. This entails that they are no longer necessary, nor arbitrary – as Kant considered them – but become particular and reliable. Indeed, whereas Kant considered the structures through which we mould reality as universal – every rational or thinking being would, in his opinion, perceive and understand reality through those structures – but arbitrary – there was no way of conceiving why we structure reality the way we do - Lorenz explains this a priori as the result of natural selection. In this regard, it is no longer necessary or universal, since it is an adaptation of a particular organism, but it also loses its arbitrary character, for it must be shaped to provide us with an accurate representation of the world.

2.1.3 The experience of the lineage

Therefore, what is a priori – prior to experience – from an ontogenetic level (the level of the organism), becomes a posteriori – posterior to experience – on a phylogenetic level (the level of the species) (Wuketits, 1990). We might look at reality through spectacles that structure it and therefore not see the world as it is in itself, but evolution by natural selection entails that the world itself shaped these spectacles, and therefore ensures that the picture of reality they provide us with is a truthful one. The particular way in which we perceive and understand reality is the result of the experience gathered by our species through evolution. Accurate ways of representing reality are kept at the expense of inaccurate ones. As Lorenz phrases it:

> The 'spectacles' of our modes of thought and perception, such as causality, substance, quality, time and place, are functions of a neuro-sensory organization that has evolved in the service of survival. When we look through these 'spectacles', therefore, we do not see, as transcendental idealists assume, some unpredictable distortion of reality which does not correspond in the least with things as they really are, and therefore cannot be regarded as an image of the outer world (Lorenz, 1977:7).

The Lorenzian argument, therefore, rejects both empiricism and transcendental idealism. As Buskes (1998) explains, the empiricist's tenet, claiming that all knowledge is derived from experience, is false, because humans – like any other organism for that matter – are endowed with innate, a priori knowledge. Kant's transcendental idealism, on the other hand, which argues that the world in itself is necessarily beyond our ken since it is moulded by the subject, is – according to the Lorenzian epistemologist – equally mistaken, because our perceptual and cognitive apparatus conform to a pre-structured world by the process of natural selection (43).

Indeed, according to Hoffmeyer (1993), with every generation, a meeting between the species and its environment occurs – each time with a different outcome – as every generation passes on a different set of genetic material to the next generation. The genes that provide an organism with the best ways of coping with the environment will be more present in the next generation than in the previous one, because the organisms endowed with these genes are more successful in reproducing. Every single generation, however, is faced with a unique set of conditions, since environments are constantly changing. This means, from the point of view of the lineage or the gene-pool of the species, that it constantly adapts in response to the experience it gathers from the confrontation of every generation with the environment (21-22).

This ongoing experience of the lineage through consecutive generations, it is argued, ensures a correspondence between cognition and nature. While we are endowed with an a priori (an innate way of perceiving and understanding the world; in other words, the 'spectacles' through which we view the world) that we cannot transcend on the ontogenetic level, the process of evolution by natural selection entails that this a priori is the product of a posteriori knowledge (i.e. knowledge gained from experience) on the phylogenetic level. The lineage, in other words, acquires knowledge of the 'world in itself' through experience, and this knowledge is reflected in the working of our perceptual and cognitive abilities. The representation of the world they provide us with must, therefore, correspond to the world itself.

In this regard, the view of the world that emerges from our perceptual and cognitive abilities can be seen, according to Vollmer (1984), as hypotheses about the structures of the real world. These hypotheses are unconscious and uncritical, and therefore utterly incorrigible. However, through the process of evolution by natural selection, (heavily) mistaken hypotheses soon die out and are replaced by better trials and better conjectures, enhancing the chances of survival and reproduction. Indeed, if these 'a priori conjectures' about the world are relevant for survival and reproduction, natural selection will gradually improve them, as it does for any other characteristic of the organism. This process of trial and error; mutation and selection, must eventually lead to a far-reaching similarity between an organism's representation of reality and the objective structures of reality (75). So, at least, goes the argumentation of many evolutionary epistemologists (see also Wuketits, 1990:88).

2.1.4 The enigmatic congruence between mind and the world

This view of natural selection ensuring the correspondence between our mental a priori and the external world, offers an explanation to a philosophical conundrum that has baffled philosophers for ages and that has led to a wide variety of theories. How is it possible that our mental structures apply to the external world? How, in other words, can external events and properties be predicted on the basis of a priori, mental models? Why is – as Galileo (1623) phrased it – the great book of nature written in the language of mathematics?

Plato (380 B.C.) thought this ability originated in the pre-existence of the soul, which contemplated the Ideas – the essence of reality – freely, before being cast in an earthly body. Descartes (1641) 'proved' the existence of a perfect God – and therefore one that would not deceive us – to explain why our innate ideas enable us to understand reality. Leibniz (1704) invoked a 'pre-established harmony' between the subject and the objective world as given by the hand of God. Kant's (1781) transcendental philosophy turns the tables on these traditional approaches, and holds that our cognitive structuring does not match the world, but rather that it is the experienced world – through the categories of perception and understanding – that matches our cognitive structuring. Our mind – in other words – is not shaped to unveil the objective structures of reality, but our subjective experience of reality is shaped in a way that allows us to acquire knowledge of it. Lorenzian evolutionary epistemology, on the other hand, explains this congruence between our cognitive structures and external reality through the process of adaptation. Our mental a priori, it claims, is the product of a process shaping it in accordance with the external world.

2.2 Hypothetical realism

The claim that the representation of the world that we acquire through our 'a priori veil' – i.e. through our perceptual and cognitive apparatus – matches the structures of the real world, can therefore, according to the Lorenzian argument, not be doubted. In the meantime – as Lorenz (1941, 1973) was fully aware – natural selection does not ensure an autonomous and absolute validity of man's grasp of reality. They do not have a priori validity; they are not necessarily true, nor infallible, but merely provide us with decent approximations of what really is the case. In other words, they provide us with a truthful 'basic' representation of the world, but not one that mirrors the world in all its inherent complexity.

This approximate character of cognition is elegantly illustrated by Lorenz in the following comparison:

> [Just as] the 'dots' produced by the coarse 'screens' used in the reproductions of photographs in our daily papers are satisfactory representations when looked at superficially, but cannot stand closer inspection with a magnifying glass, [so], too, the reproductions of the world by our forms of intuition and categories break down as soon as they are required to give a somewhat closer representation of their objects (Lorenz, 1941:128).

Therefore, it is argued, evolutionary epistemology leads to 'hypothetical realism' (Vollmer, 1984:78). It concludes, from the fact that our cognitive system is shaped by natural selection, that there is a far-reaching agreement – albeit an incomplete and imperfect approximation – between the objective structures of the world and our mental reconstruction of those structures through perception and cognition.

Vollmer (1984) points out the impossibility of a perfect agreement. Such an ideal adaptation, he claims, is not necessary for survival, and could only occur at a very high cost. Evolutionary change does not come for free. Homo sapiens' increased brain power necessitated an equally increased amount of nutrients to keep the organism going. Therefore, there is always a trade-off between advantages procured by a new ability or physical trait and the cost of this adaptation. Furthermore, he claims, a perfect fit would mean a very rigid 'reflection' of nature, leaving no chance of survival in the case of environmental changes (78).

This hypothetical realism forms the backbone of Lorenzian evolutionary epistemology. Admitting that our representation of the world might never match the world itself perfectly, it argues that it must nevertheless be a decent approximation of it (see Lorenz 1941, 1971; Vollmer 1975, 1984; Riedl 1984b; Oeser 1987; Wuketits 1990). This train of thought is at the core of Quine's famous quote with respect to the infamous problem of induction.

> If people's innate spacing of qualities is a gene-linked trait, then the spacing that has made for the most successful inductions will have tended to predominate through natural selection. Creatures inveterately wrong in their inductions have a pathetic but praiseworthy tendency to die before reproducing their kind (Quine, 1969:126).

At first glance, within the framework of evolution by natural selection, this claim seems very plausible. How, indeed, could a more accurate representation of the world - a more correct way of inducing - not profit an organism in terms of survival and reproduction? However, as I will argue, this view commits the fallacy of deducing the truthfulness of our representation of the world from its mere functionality with regards to biological fitness. In order to show this, I will first point out that natural selection, only shaping perceptual and cognitive abilities to produce desired behaviour (cf. previous chapter), cannot be expected to yield true beliefs about the world automatically. In doing so, I will give an overview of the possible relations between beliefs and behaviour, based on Plantinga's (1993) argument against evolutionary naturalism. Moreover, I will argue, the assumption that true inferences have more survival and reproductive value than false ones is not as self-evident as it appears. Furthermore, a closer look at variation as the motor of adaptation, and the importance of both the particular elements of the environment and previously acquired traits in the selection of new ones, points to the constitution of an imperfect and very particular, species-specific way of representing reality.

3. Uncovering the fallacy

3.1 Plantinga's case against evolutionary naturalism

Alvin Plantinga (1993, 1994) caused quite a stir with an analytical argument claiming that evolutionary naturalism is self-defeating. At first glance, this seems an outrageous and even absurd claim. Indeed, it is precisely the theory of evolution that provides a naturalistic account of the organic world, its apparent design, and even the appearance of mind. As Richard Dawkins (1986:6-7), the famous British evolutionary scientist, points out, atheism might have been logically tenable in the pre-Darwinian era, but it is nevertheless Darwin's theory of evolution by natural selection that made it possible to be 'an intellectually fulfilled atheist'. Indeed, how could one have explained the obvious order and design in nature, the fact that all organisms are interdependent on each other and that one species – namely our own – is endowed with conscious thought, before Darwinism entered the stage? How could this undeniable design in nature be the product of mindless causality in a purely material universe?

Plantinga (1994), however, detects in this inseparable pair, i.e. evolution and naturalism, a core of contradiction. Naturalism – meaning that there is no supernatural entity involved in the world – and evolution, he claims, are at odds with each other. Indeed, if both are true, the probability that our cognitive faculties are reliable is low or at least inscrutable – meaning that we cannot assess their reliability (1). For Plantinga, cognitive faculties are reliable if 'the great bulk' of its deliverances are true. If naturalism is true, there is no God or other supernatural entity creating mankind in its image and therefore ensuring us that our reasoning faculties are reliable. Evolutionary naturalism entails that all organisms, including mankind and its cognitive faculties, are the result of selected random variation (2-3). The problem that brings us to the core of the argument is that – as pointed out – natural selection does not care about beliefs, only about behaviour. Believing something does not enhance survival and reproductive success; only particular behaviour does. Therefore in order for natural selection to have 'a grip' on cognitive faculties, they have to be causally related to behaviour.

3.1.1 Possible relations between beliefs and behaviour

Plantinga (1994) goes on to enumerate the possible relations between beliefs and behaviour. The first possibility is known as 'epiphenomenalism'. Under this view, beliefs cannot exert any influence on the material world. They are in no way connected with behaviour, since only matter can change matter, not mind. This is one way of solving the infamous body-mind problem which asks how mind could be connected with matter. It claims that they are two completely different and separate substances. There is no interaction whatsoever. On this perspective, of course, the probability that our cognitive faculties are reliable is low, since natural selection won't have any grip on them (6-7).

The second possibility is known as 'semantic epiphenomenalism'. Beliefs do have a causative link to behaviour. However, they cause behaviour not in virtue of their semantic properties – their content or meaning – but only in virtue of their syntactic properties. Behaviour is influenced by electrochemical properties in the brain, not by the content these electrochemical properties cause us to believe. The proposition itself - the statement which can either be true or false - remains, in this view, out of the grip of evolution. Therefore, once again, the probability is low that the beliefs produced by our cognitive faculties are reliable (7-9).

The third possibility is that our beliefs cause behaviour – both semantically and syntactically – so the content of our beliefs influence behaviour, but these beliefs are maladaptive. Organisms would be better off without them. Of course, in this case too, the probability that beliefs are reliable is low, since they actually prevent appropriate behaviour and natural selection would strive to eliminate them (9).

The final possibility is the one adhered to in our commonsense thinking about the relation between behaviour and beliefs or body and mind, and is the one underlying the Lorenzian school of thought which links cognition and reality. It states that our beliefs cause behaviour, semantically and syntactically, and that these beliefs are adaptive. In other words, our beliefs lead to appropriate behaviour. In this case, natural selection does have a grip on beliefs. It selects those that lead to appropriate behaviour and eliminates those that do not (10). In this regard, one is tempted to conclude – as do the above mentioned evolutionary epistemologists – that our beliefs must be, at least to some extent, reliable, since it seems almost indubitable that the more reliable a belief is, the more appropriate the resulting behaviour will be. How, indeed, could mistaken beliefs about the world lead to better behaviour in the world?

3.1.2 Belief-cum-desire systems

However, as Plantinga (1993) argues, the probability that these beliefs would be accurate isn't nearly as high as one would be inclined to think. The reason he offers in support of this is that behaviour is not only caused by belief but also by desire. In this light, any kind of adaptive action could be caused by a wide variety of 'belief-cum-desire' combinations. To illustrate this, he gives the example of a prehistoric hominid, Paul, running away from a hungry tiger. This appropriate action, however, can be caused by a lot of different belief-desire combinations, leaving no guarantee as to the reliability of the isolated beliefs.

> Perhaps Paul very much likes the idea of being eaten, but when he sees a tiger, always runs off looking for a better prospect, because he thinks it unlikely the tiger he sees will eat him. This will get his body parts in the right place so far as survival is concerned, without involving much by way of true belief. … Or perhaps he thinks the tiger is a large, friendly, cuddly pussy cat and wants to pet it; but he also thinks the best way to pet it is to run away from it. … or perhaps he thinks the tiger is a regularly recurring illusion, and, hoping to keep his weight down, has formed the resolution to run a mile at top speed whenever presented with such an illusion; or perhaps he thinks he is about to take part in a 1600m race, wants to win, and believes the appearance of the tiger is the starting signal; or perhaps… . Clearly there are any number of belief-cum-desire systems that equally fit a given bit of behavior (Platinga, 1993:225-226).

These examples, of course, strike us as very implausible. The point they are trying to make, however, is right on the mark. Since natural selection only shapes perceptual and cognitive abilities to yield particular behavioural schemes that favour an organism's chances for successful reproduction, we cannot assume that this process provides us with accurate beliefs. As shown in the previous chapter, flawed beliefs in folk physics or biology can provide a suitable basis for behavioural interaction with the environment that is highly adaptive.

3.1.3 Conclusion

Plantinga concludes, therefore, that since these four possibilities are jointly exhaustive, and since all lead to the conclusion that, reasoning from an evolutionary naturalistic perspective, the reliability of our beliefs are low or at least inscrutable, evolutionary naturalism is self-defeating. Indeed, it is itself a belief, and thus the probability of it being true is low or inscrutable, leading to the conclusion that, from an evolutionary naturalistic perspective, evolutionary naturalism is improbable. To this perceived self-defeating aspect of evolutionary naturalism later, however, I will return later (cf. Chapter 6).

At this point, Plantinga's argument is offered in support of the claim that evolution in no way cares for the truth of the interpretations held by the organism it shapes. There is absolutely no way of deriving the truth or accuracy of our beliefs about reality from the fact that they are the product of abilities shaped by natural selection, precisely because natural selection only shapes cognition to yield appropriate behaviour.[11] In short, the mindless process of evolution by means of natural selection is driven by survival and reproduction, and these cannot account for the truthfulness of our representation of the world. In some cases, as I will argue below, it can even be shown to prevent rational or true beliefs.

---

[11] A similar argument is held by the neurophysiologist and philosopher of mind, Patricia Churchland (1987). She claims that the main function of human brains is to allow the organisms to move in appropriate ways – this means to make the organism succeed in 'the four F's': feeding, fleeing, fighting and reproducing. Another way of representing the world will only be selected for as long as 'it is geared' to the organism's biological needs, increasing its chances for survival and reproduction (548).

3.2 Natural selection and rationality

3.2.1 Lorenzian evolutionary epistemology's problematic assumptions

The philosopher Stephen Stich (1990), points out that the common argument of (Lorenzian) evolutionary epistemology – i.e. that natural selection guarantees the accuracy of our beliefs or the rationality of our thinking – is mistaken. There are, indeed, two unproven assumptions in this claim - firstly, that evolution produces optimally well-designed systems, and secondly, that an optimally well-designed system is a rational one. Both statements appear problematic under closer examination (56).

The first statement claims that since the process of natural selection selects the genes endowing the organism with the highest chance for survival and reproduction, it will eventually yield the best possible cognitive system for that organism - the best in terms of biological fitness, that is. In this light, one can say that a cognitive system is better than another when it enhances an organism's chances for survival and reproduction to a greater extent.

The second statement claims that such an optimally well-designed system must be a rational one, meaning that it produces good reasoning or just inferences, yielding a true representation of the world. This is the main argument of the evolutionary epistemologists cited above. A system that is biologically fit must be a rational cognitive system and produce at least an approximately accurate representation of what really is the case. The more a system is fitness enhancing, they argue, the more rational it will be. Why? Because having true beliefs rather than false ones increases an organism's chance of survival and reproduction.

3.2.2 Does evolution produce optimally well designed systems?

With respect to the first statement, things are not as straightforward as they appear. It can be said that a cognitive system is better than another one when the 'input/output pairings' it provides enhances an organism's chances for reproductive success (Stich, 1990:60). However, there is another factor to take into consideration. A cognitive system will also be better than another one if it offers the same biological fitness to the organism at a lower cost. This is pointed out by Elliott Sober (1981). There are both external and internal considerations to take into account in order to assess biological fitness.

External considerations look at the effectiveness in achieving the intended purpose (in this case, survival and reproduction) and internal considerations look at the costs - the amount of resources the system requires (in this case the demands on the memory, energy and other resources of the organism). At some point, the benefit of additional information extracted from the environment will not weigh up against the additional resources this requires from the organism. Fitness, in this regard, is not just a function of the abilities of an organism, but also of the demands these abilities make.

Therefore, the best possible cognitive system will not automatically be selected. There is a trade-off between external fitness of the system – how well the abilities enhance survival and reproduction in a given environment – and internal fitness – how much these abilities cost the organism or how much resources have to be committed to the particular abilities which therefore cannot be spent on other needs. This argument, however, is accepted by most Lorenzian evolutionary epistemologists and is often mentioned in order to point out a 'hypothetical realism' (cf. Vollmer, 1984). The second assumption, on the other hand, is non-negotiable for Lorenzian epistemologists. The process of natural selection shapes our cognitive system to yield a true representation of the world. Stich will point out cases where this statement fails.

3.2.3 Are optimally well designed systems rational ones?

Stich (1990) claims that the best possible system – in terms of biological fitness – is not automatically the most rational one – i.e. the most reliable inferential system. Even if one system is less reliable than another one – i.e. in that it produces less accurate beliefs – it is nevertheless possible that it will exceed the more reliable system not only in terms of internal fitness – in that it could run more economically – but also in external fitness, and, therefore, will obviously be favoured by natural selection.

Stich starts by distinguishing between two different kinds of wrong inferences. The first one is to infer that something is the case, when in fact it is not. The second one is the opposite, inferring that something is not the case, when in fact it is. The former are referred to as 'false positives', while the latter are 'false negatives' (61). A false positive, for instance, would be to think a plant is poisonous when in fact it is not. A false negative, on the other hand, would be to think a plant is not poisonous when in fact it is. These strategically chosen examples bring us to the core of the argument. Some mistakes are cheap, while others are very expensive. In this case, the false positive would be relatively cheap – missed nutrition which has only a small negative impact on the general chances for survival and reproduction of the organism – while the false negative would be very expensive, intoxicating the organism and possibly ending its life.

Therefore, as Stich points out, an inferential strategy bringing about a high level of external fitness would be very risk-averse when assessing which plants are poisonous and which are not. Such a conservative strategy would undoubtedly generate an important number of false positives, since it would consider a plant poisonous on relatively weak evidence (e.g. vague resemblance to a known poisonous plant, or even an assumption that it is poisonous based on a lack of information). Since false negatives, in this case, are lethal, and false positives are relatively 'cheap' mistakes, the inferential system leading to the highest level of external fitness would be one that makes a considerable number of cheap mistakes, thus avoiding the expensive ones. In this light, a system might provide more fitness to an organism and still make more mistakes, therefore being less reliable. The most reliable system would indeed be bound to make more expensive mistakes, since it would only strive to reduce the absolute number of mistakes made, without distinguishing between kinds of mistakes and their respective consequences. Natural selection, therefore, will in some cases provide us with error-prone systems due to conservative strategies – yielding a substantial amount of false positives as in the aforementioned example – rather than more reliable systems – yielding far fewer false positives at the cost of a few more false negatives (62).

Of course, as Stich is fully aware, the ideal inferential strategy would be one that always yields the right answer. However, as we pointed out above, such a system – even if it were available – would be very expensive in terms of internal fitness. Indeed, when the utility of further information does not exceed the cost which this additional ability entails for the organism in terms of energy, time, and so on, it will not be selected. Furthermore, it seems very dubitable to even suppose that this could in theory be possible, as the environment is continually subject to change. Other organisms evolve new ways of deception (i.e. edible plants evolve characteristics resembling poisonous ones to avoid being eaten), predators evolve new camouflages, parasites new resistances, and so on. Therefore, danger detection is bound to be imperfect and bound to make a trade-off between 'overall reliability' and 'reliability when it matters' (62-63).

Since natural selection does not care in the slightest for truth, but only for reproductive success, it will not shape our cognitive system to provide us with the best possible grasp of what really is the case, only with inferences that – even when they're often mistaken – guide our actions in such a way so as to maximise our chances for survival and reproduction.

3.2.4 Cognitive biases and error management theory

In this context, studies in cognitive sciences have brought to light a myriad of cognitive biases that human thinking is subject to. (Kahneman, Tversky, 1973; Kahneman, Slovic, Tversky, 1982; Piattelli-Palmarini, 1994; Sutherland, 1994). These 'inevitable illusions', as Piattelli-Palmarini (1994) calls them, are the product of 'heuristics', or mental strategies for information processing, that in some cases lead to erroneous inferences. Just as in the case of visual illusions, we cannot help but fall into the trap of this predisposition towards irrational inferences in certain particular situations. In this sense, cognitive illusions are general, in that all human beings are subjected to them; systematic, in that they can be reproduced in different contexts; directional, in that their effects tend in one direction (i.e. rounding off, framing, segregating, and so on); and, as pointed out above, incorrigible by logic, in that even when we see the error, we are still spontaneously inclined to make it again. (Piattelli-Palmarini, 1994:139-140). In the same way as our vision – most often reliable and accurate – can fool us in particular circumstances, so can our cognitive abilities fool us under specific conditions. They evolved in order to provide us with accurate inferences in most cases, but in particular conditions they lead us systematically to mistaken inferences.

A typical example of such an illusion is known as 'the gambler fallacy'. When playing roulette, after ten consecutive reds, people will be very tempted to bet on 'black', feeling they're 'overdue for black'. Of course, statistically, chances remain 50/50. Another one is known as probability blindness. Apparently, we're disposed to pay more for something that reduces the risk from one in a thousand to zero, than from two in a thousand to one in a thousand, or from four to three in a thousand. Another example of such 'blindness' is the difference in our perceived incentive to purchase extra lottery tickets, for instance, to increase the probability of winning, according to the range in which the probability hovers. We'll be much less willing to buy extra tickets when our chances go up to 37% from 32%, than in the case where we can boost our chances to 99% from 94%. Mathematically, of course, the cost-benefit calculation should be the same in both cases: a 5% gain in probability. (Piatelli-Palmarini, 1994:130-131).[12]

The cause of these biases, as Stich (1990) predicts, relates to both internal fitness - weighing general results against the cost supported by the system and endowing us with inferences that work in most cases, but which can break down in (atypical) particular circumstances – and external fitness – originating in the selective preference for biased reasoning, which leads to more low-cost errors, rather than unbiased reasoning, which reduces the total amount of errors, but results in more high-cost errors.

This last consideration gave rise to what the psychologists Haselton and Buss (2000) labelled 'Error Management Theory'. Error management theory is a theory about judgements made under conditions of uncertainty. It predicts that if false positives and false negatives entailed different costs over the course of evolutionary history, natural selection will have biased our inferential system in order to reduce the amount of potential costs, regardless of the absolute number of mistakes this causes. Building on this insight, Haselton (2006), together with a fellow psychologist, Nettle, argues that the human psyche evolved into a 'paranoid optimistic' decision maker. Depending on the context – associated with a particular pattern of potential

[12] With respect to intuitively biased probalistic calculations, Gigerenzer and his colleagues point out that these apply more to a 'Baysian' than a 'frequentist' take on probability. Whereas Baysians view probability as the subjective degree of confidence one can have in predicting, frequentists view it as the frequency with which events occur in the world. When problems are formulated in terms of the latter framework, intuition seems to serve us better (Gigerenzer, 1991). Tooby and Cosmides point out that this can be explained in an evolutionary context, where our ancestors only encountered problems of the second kind. (Cosmides & Tooby, 1996).

benefits versus costs throughout evolutionary history – we either take a very conservative (paranoid) stance, and are not willing to take any chances, or, on the contrary, we appear to throw caution to the wind, 'blinded' by the anticipation of success. Cognitive biases, in this light, incite us to adopt a belief in one domain on much sparser evidence than a belief in another domain.

It has been shown, for instance, that we systematically underestimate the time-to-impact of approaching sounds, i.e. we always expect the object to arrive faster than it really does (Neuhoff, 1998). On the other hand, men apparently, have a tendency to overestimate women's sexual intent towards them (Haselton & Buss, 2000). Both biases make sense in an evolutionary perspective. Indeed, with regard to the former, the additional cost of expecting the impact too soon is very low, whereas the risk of expecting it just a fraction of a second too late is, of course, rather high. Therefore, natural selection biased us to expect it too soon, leaving us prepared for what's coming. As for men's sexual 'over perception', it is similarly obvious that the costs of missing a possibility of reproduction is very high in terms of biological fitness – which is directly linked to reproductive success – whereas the cost of overestimating the possibility of a sexual encounter rarely exceeds a slap on the cheek.

3.2.5 Non-cognitive purposes

Tackling the question as to whether natural selection leads to rational thinking from a different perspective, Peter Munz (1993) points out that theories often fulfil functions other than representing the environment, stemming out of the obvious evolutionary incentive to extract useful elements and avoid harmful ones. Indeed, shared beliefs seem to play a crucial role in what Munz labels 'non cognitive purposes', such as, for instance, social bonding. False beliefs, in this context, will thrive precisely in virtue of this non-cognitive purpose. Indeed, because of their particular content which is adhered to by a particular group, these beliefs operate to set out boundaries for a social group, including believers and excluding non-believers. In order to achieve this purpose, these beliefs cannot be obviously accurate or truthful – rational – beliefs, because such beliefs would be accepted by any sound minded person, which, of course, defeats the whole purpose (169-170).

In this regard, it can be argued that natural selection has provided us with a cognitive system that not only deals with reality in a representative way – a way that incites us to act appropriately – but that also forms irrational beliefs, satisfying other purposes. This aspect of human belief-forming is at the core of the wide variety of mythological, religious and other esoteric accounts of reality. While in the age of enlightenment such theories were often seen as primitive and superfluous in our scientific era, and as soon to be eradicated, it appears that they are not simply replaceable by rational, scientific thought. This is precisely because they stem from a different aspect of our cognitive system, adapted to fulfil other, non-cognitive purposes, such as providing social boundaries. The importance of such beliefs in delimiting social groups is, indeed, rather obvious throughout history, which has been infested by wars and conflicts based on religious and cultural differences. It is precisely this aspect of humanity that made Schopenhauer (1819) characterise man as 'homo metaphysicus'. We do not limit the representation of our world to objective, realist descriptions, but seem to have a constant urge to go beyond what is manifestly given.

3.2.6 Conclusion

We are not 'programmed' to develop expectations that reflect what really is the case as accurately as possible, but to develop expectations that direct our actions in such a way so as to maximise our biological fitness. As Plantinga (1993) rightly claimed, natural selection does only have a grip on behaviour. Cognition, as in expectation and belief-forming, is merely there to enable us to act in appropriate ways. Therefore, the central argument of evolutionary epistemology – that natural selection shaped our perceptual and cognitive abilities to yield an accurate representation of the world – fails.

In this context, Stich (1990) shows that natural selection leads not only to less than optimally well designed systems (internal fitness), but, moreover, that optimally well designed systems in terms of fitness are not per se the most rational ones. Indeed, some cognitive schemata may generate a larger number of false beliefs than other schemata, and nevertheless endow us with 'fitter' behaviour than the latter. When this is the case, these will be selected and leave us with biases, systematically deceiving us in a particular way. Finally, Munz (1993) points at non-cognitive purposes that our belief-forming system serves. Precisely because of their irrationality, these beliefs will be adopted. Therefore, arguing from an evolutionary perspective, we cannot infer the rationality or accuracy of our representation of the world from the fact that this representation is the product of abilities evolved by natural selection.

3.3 The mechanism of evolution

Besides these compelling arguments showing that biological fitness does not automatically imply truthful representations, a closer look at the mechanism of evolution itself provides further reason to reject Lorenz's claim. Indeed, as I will show below, the human mind, being the outcome of an evolutionary process in the same way as any other adaptation, must be both an imperfect and utterly particular organ.

3.3.1 Undirected variation

Evolution by natural selection happens because of the selective retention of genetic variation endowing organisms with phenotypic advantages (i.e. characteristics enhancing its chances for reproduction). But what causes variation? Darwin himself could not provide an answer to that question. However, when Mendelian genetics – uncovering the basic mechanisms of inheritance – was combined with Darwinism, in what would become known as 'the synthetic theory of evolution' or 'the modern synthesis', the chips fell into place.

The two main causes of variation are genetic mutation and genetic recombination.[13] The first cause, mutation, happens in the case of a replication error in cell division. As Mayr (2001) points out, although the replication of DNA molecules during cell division and gamete formation is very accurate, the occasional error can occur. When this happens, DNA is altered and subsequently inherited, until a new error arises. The mutation can either be beneficial, neutral or deleterious to the chances of reproduction of the organism. When it is beneficial – which again is very rare – it will be favoured by natural selection and spread over the gene-

[13] A third source of genetic variation is the rare case of 'unidirectional lateral transfer' found in asexually reproducing bacteria, where a bacterium attaches itself to another bacterium and transfers some of its genes. Interestingly enough, this process apparently found its way into sexual reproduction, where one gene can move from their position on one chromosome to another chromosome. Whether this phenomenon is common, however, is not yet known. (Mayr 2001:117).

pool. This is the way in which genetic variation arises in asexual reproduction, where the offspring are clones of the reproducing organism (106-108).

With sexual reproduction, however, a new source of genetic variation enters the scene. In this case, the genes of both parent organisms are recombined in the offspring in a random fashion. The result is a unique genotype, different from either parent organism. This way of reproducing speeds up the process of evolution in a considerable way, providing natural selection with a far greater number of new phenotypes to select from (114-115). However, as Dercole and Rinaldi (2008:9) point out, with only genetic recombination, where the genetic material is a mix of the parents genes, no new phenotypic value – meaning increased adaptation of a phenotype to a constant environment – can appear, other than those already observed or potentially observable by suitable mixing of individual genotypes. In order to alter the gene-pool of a species, therefore, mutations are necessary. Mutations, in this regard, are the ultimate source of genetic variation.

What matters here, however, is the randomness of genetic variation. The production of variation through mutation and recombination is completely and utterly undirected.[14] Variation arises independently of what is needed for adaptation. In this light, Mayr (2001:131-132) characterises natural selection as a two-step process. It is both a process of chance or accident (the production of genetic variation) and determinism (the testing of this variation, resulting in the selective retention of beneficial variations to the detriment of others).

Ridley (1993) argues that there are both factual and theoretical reasons to suppose variation is undirected. The factual evidence comes from laboratory observation of spontaneous mutations. These do not arise in relation to the adaptive needs of the organisms. The theoretical reason, on the other hand, points at the inexistence of a genetic mechanism that could direct the right changes to happen. Such a mechanism, Ridley points out, would be practically impossible. The organism would have to recognise that the environment has changed, infer what change is needed to adapt to these new conditions and then cause this

[14] Ridley (1993:79) points out, however, that whereas variation is random with respect to the direction of adaptation – meaning that the necessary adaptation can only be generated by chance – it does not exclude the possibility that mutations are non-random at the molecular level. When I characterise variation as random or undirected, I do so in relation to adaptation.

change to happen in the relevant parts of its DNA. In other words, they would have to describe a subject matter they have never encountered (a new environment) in a language they do not understand (the information encoded in DNA), '[l]ike a seventeenth century American using Egyptian hieroglyphics to describe how to change a computer program' (Ridley, 1993:78).

3.3.2 Evolution as a gradual, non-teleological process

As Mayr (2001) emphasises, natural selection is not goal-directed. It is a mindless or blind process without any form of long-term vision. It simply eliminates, repeatedly in each generation, genes that have proved less adaptive than others in their phenotypic effect. It does not, in other words, strive to gradually alter a simple bacterium into a complex, intelligent creature or to systematically improve the adaptations in a particular direction. Lineages often end in extinction or are driven in a radically different direction in response to a new challenge from the environment (133).

There is, in this regard, absolutely no way of predicting how species will evolve. Indeed, different genotypes within a species may respond differently to the same change of environment, making it unpredictable which genotypes are going to be selected. Furthermore, the changes in the environment, pushing the organisms to adapt, are completely unpredictable, and finally, there are no known genetic mechanisms that could produce teleological processes (133-134).

Moreover, evolution is gradual. This means that every single genotypic alteration of the phenotype has to procure an advantage to the organism in terms of reproductive fitness. A species can only acquire new characteristics and abilities as long as every little step in their constitution provides it with an advantage. In this sense, qualities that would be adaptive, but that require mutations that do not have a direct and immediate beneficial effect, cannot be selected, and neither can mutations that would be adaptive in combination with other mutations, but not on their own. Natural selection, in this sense, is an utterly myopic architect, only looking one step ahead.

All organisms, therefore, are contingent products of a blind process. They are the result of an endless chain of coincidences. Both the appearance of particular variation and the selective forces retaining the most fit alternatives in the light of particular conditions, are responsible for the evolution of a species. The evolutionary road that species take, in this regard, is the result of a completely and utterly undirected process. This process has no long term vision, but shapes organisms in an unpredictable, non-teleological and, therefore, imperfect way.

3.3.3 Path dependency

Adaptation is problem-solving, not engineering. When species develop new adaptations, these do not arise out of nowhere, but rather alter the existing model as a result of new requirements. All organisms have unnecessary, non-adaptive traits, such as, for instance, the femur bone of the whale or the human coccyx. These are leftovers from the evolutionary history of the species and have not been cancelled out completely, because their presence is as good as neutral with regard to fitness. In this context, evolutionary biologists do not talk about intelligent, but rather about stupid design. Buskes (2006) illustrates this imperfect design with the vertebrate eye. Often cited as the most obvious example of complex design, it appears, however, under closer examination, to have a remarkable flaw. This is that the nerve fibres do not connect at the back of the eye (as is the case for the eye of the squid, for instance), but in the front, and from there move back again, through the retina, to the brain. Because of this design, there is a tiny hole in the retina, causing a blind spot in each eye. This problem is then solved by natural selection by combining the vision of both eyes, each with a different blind spot, which can therefore be neutralised in united vision (282). According to Badcock (2000), this imperfect design results from the fact that the light-sensitive skin cells that gradually evolved into the vertebrate retina were under the skin, rather than on top of it (19). This is a perfect illustration of natural selection at work. Instead of optimally designing the eye for its purpose, it can only solve the problems created by a less than perfect design. Because it has no foresight, it cannot engineer a perfect organ, but can only adapt imperfect organs to particular problems.

The reason for this is that evolution is path-dependent. It depends on what has already been acquired. New adaptations are selected not only in the light of the purpose they are selected for, but also in the light of the basic design already in place. Evolution cannot erase parts of the existing blue-print in creating a new adaptation, but merely transforms the blue-print to fit new requirements as well as possible. In this context, the human dorsal column is apparently not as well suited for an upright position as it was for our four legged ancestors. When we starting walking on our two feet, however, natural selection could not rethink the whole structure, but merely adapted the existing structure to the new requirements, leaving quite a few people with a sore back (Buskes, 2006:282).

Path-dependency, therefore, is responsible for both sub-optimal design – as in the case of the vertebrate eye or the human dorsal column – and for the restriction of possible adaptations in the light of previous ones, resulting in characteristics and abilities that are continuous with the evolutionary history of the species. Indeed, every acquired characteristic or ability of any organism throughout evolutionary history bears the marks of past acquisitions, all the way down to the very structure of our cells. In this regard, our perceptual and cognitive abilities can be seen as particular 'enhancements' based on a primate model, which again is grounded in a mammalian and then a vertebrate model, all the way down to the first prokaryotic organisms. This yields a particular and imperfect way of adapting to the environment. Indeed, as pointed out, a new environmental challenge will always be met by altering some of the traits and abilities which are already in place, not by starting from zero and selecting variations in the creation of an optimal new adaptation. Since it is often necessary to take a step back in order to take three forward, and natural selection cannot backtrack, species can and do get stuck in sub-optima.

3.3.4 The environment

To what is an organism adapted? The common answer is that organisms are adapted to their environment. However, as pointed out in the previous chapter (cf. Chapter 1, 4.1.1 Detecting 'affordances'), species are adapted only to particular properties of the environment. Two different organisms living in the same environment, therefore, can be adapted to very different aspects of that environment. The properties of the environment relevant for a particular organism are what Gibson (1979) calls 'affordances'. The specific set of affordances for a particular organism constitutes its niche. Adaptations, through the working of natural selection, are therefore shaped to answer the challenges presented by this particular niche, with both its beneficial and harmful affordances. In this context, homo sapiens too – including its perceptual and cognitive abilities – is adapted to particular aspects of a particular environment: its very own niche.[15]

Building on this insight, Gerhard Vollmer (1975, 1984) argues that we are, in the same way as any other organisms, endowed with a 'cognitive niche' – a part of the world with which we coped in our evolutionary history and to which our perceptual and cognitive apparatus is adapted. This niche in which we find ourselves, or this 'umwelt', to use Uexküll's (1909) term, is referred to by Vollmer as 'mesocosm'. Evolutionary psychologists, on the other hand, argue that we are endowed with a 'stone-age mind', often very poorly adapted to contemporary, modern living conditions. They claim that much of human behavioural dispositions and intuitions about the world are the outcome of psychological and cognitive mechanisms adapted to our ancestor's way of life. (Barkow, Tooby & Cosmides, 1992; Pinker, 1997, 2002).

A) Vollmer's mesocosm

Vollmer (1975, 1984) states that human sensory and cognitive abilities are adapted to cope with a particular section of the external world: the world of medium dimensions. This part of the world, which he calls 'mesocosm', is best described as the realm of dimensions – relative to physical quantities such as spatial extension, time and mass – that are relevant in our interaction with the environment. For instance, we can 'visualise' or gauge stretches of time reaching from a mere second to decennia, distances from millimetres to kilometres or mass from grams to tons, but beyond these dimensions our intuition fails us. (Vollmer, 1984:88-89). Both the microscopic and the macroscopic world completely elude us on an intuitive level. We have, indeed, been adapted to the particular realm of dimensions in which our organism – in virtue of its size, speed, lifespan and environment – is encapsulated. Besides

[15] This, as Gibson (1979) is fully aware, is not a denial that all species on our planet share a set of affordances – such as, for instance, oxygen, water, gravity or light – which explains the similarity of sensory and other adaptations widely diverging species often share. The point I want to make here, is only that an environment does not unequivocally shape organisms in a uniform way, but merely provides one side of a relationship that constitutes a niche, and can therefore be met in very different ways.

micro- and macroscopic dimensions, our mesocosm also excludes aspects of the world that do not carry significant information for us, such as, for instance, magnetic fields, ultra-violet wavelengths and other potential stimuli for which we haven't evolved sensory receptors.

While Vollmer argues that natural selection has endowed us with an intuitive grasp of the world that provides us with an accurate representation only within the mesocosm, I would, however, radicalise this statement. Natural selection has not shaped our cognitive apparatus to yield an accurate representation of the environment – even within the realm of meaningful dimensions – but to provide us with a framework that enables us to interact with our mesocosm in appropriate ways. Indeed, as shown in the previous chapter, the scientifically mistaken assumptions about the world to which we intuitively adhere in our folk physics and biology can nevertheless allow us to interact with our environment in very efficient ways with regard to biological fitness (cf. Chapter 1 – 3.2.3 Species-specificity).

In this regard, even the most basic assumptions about the planet we live on can be false, but nevertheless harmless when it comes to surviving and reproducing. Indeed, the radically mistaken view that the earth is flat or that the sun orbits around the earth, commonly accepted in the pre-scientific era, did not impair our ancestors' reproductive chances, even in the slightest. While it is important for our perceptual and cognitive apparatuses to enable us to predict the probable outcomes and properties of the relevant elements in our environment, this does not entail the accuracy of the framework of underlying assumptions about the world in which these predictions are grounded. A flat world or a geocentric universe is perfectly adequate to control our mesocosm, as is an 'impetus theory' or an essentialist conception of the natural world. False premises can indeed produce accurate predictions in particular circumstances. In this sense, Aristotelian physics can function as a suitable framework to explain physical happenings on this planet, as can Newtonian physics. Both, however, have been proved wrong by Einstein's theory of relativity. Nevertheless, assuming that the latter provides the correct framework – or at least the best approximation – and that the former two are mistaken, all of them could, however, predict equally well the movement of a stone cast from a cliff.

B) A stone age mind

Pinker (2002), on the other hand, points out that some of our current ordeals are the result of a 'mismatch' between the source of our passions – which are adapted to our ancestral way of life – and modern day life. Our impulsive gorging of fatty food in anticipation of a possible famine, the male desire to engage in casual sexual encounters to maximise its chances for passing down its genes, or the shot of adrenalin whenever we are confronted with a stressful situation, are just a few examples (219). Nowadays, the outcomes of these impulses are not as beneficial as they used to be, to say the least. They saddle us with obesity and diabetes, unwanted children and abortions, and stress-related problems.

Pinker argues that this mismatch does not only apply to our emotions, but also to our intellect. Our cognitive faculties were designed to meet the needs of our stone-age ancestors, and are specialised in dealing with problems humans were faced with over the course of their evolutionary history. Our cognitive apparatus carves up reality in useful ways, and has intuitive theories predicting probable outcomes, framing the world in ways that permitted pre-historical human tribes to thrive in their natural environment.

Natural selection, therefore, cannot be said to model our cognitive apparatus in such a way that matches our representations to the world itself, but merely to particular elements of the world that were relevant to the survival and reproduction of human tribes in the Pleistocene era. As Pinker (2002) states, '[our] ways of knowing and core intuitions are suitable for the lifestyle of small groups of illiterate, stateless people who live off the land, survive by their wits, and depend on what they can carry' (221). It is tailored, in other words, to the requirements of a very particular niche that we occupied for most of our evolutionary history, not to yield a general, comprehensive understanding of the world. In this context, Tooby and Cosmides (1992, 2003) point out that our cognitive apparatus evolved to solve particular adaptive problems. These problems are the result of both the environment and the particular way our ancestors met this environment to ensure their survival and reproduction (cf. Chapter 1 – 4.2.2 Adaptive problems). This left us with a very particular way of viewing the world, since evolution provided us with an amalgam of innate and thoroughly species-specific cognitive predispositions as the outcome of a set of survival and reproduction related problems the environment held for hunter-gathering hominid tribes.

3.3.5 Genetic drift and genetic hitch-hiking

Furthermore, natural selection is not the only architect in town. Genetic drift changes the gene-pool of a species in a completely random fashion. In other words, unlike natural selection, it acts upon the genotypic frequencies in a particular population without any regard for their phenotypic effect. Indeed, because of this random drift, gene variants may disappear, not because they proved less adapted, but because they were lost in the random shuffle of genetic recombination. This is referred to as 'the sampling problem'. Imagine a population of 10 individuals: 5 have the genetic variant A and 5 have the genetic variant B. Both genes are different but more or less similar in terms of fitness. The offspring would most likely inherit both variants in equal amounts. However, since this new generation's genetic constitution is a random sampling from the parental generation's, it could well be that, for instance, 70% end up with A and only 30% with B, just as it would be possible to end up with seven heads and three tails after tossing a coin 10 times. This could end in the fixation of genetic variant A and the loss of variant B. Of course, the larger a population, the less chance there is that genetic variants or 'alleles' are going to drift out of existence by this process. If you flip a coin 4 times, your chances of ending up with 3 heads and 1 tail are much higher than of ending up with 3000 heads and 1000 tails after flipping it 4000 times (Ridley, 1993:138-140).

Genetic hitch-hiking, on the other hand, happens when genetic changes at one locus entail related changes at linked loci, because of their proximity on the chromosome. The linked locus is, in other words, attached to the first locus, and is therefore unlikely to be detached in the process of genetic recombination from parent to offspring. Now, suppose natural selection favours allele A over B at one locus, but A is linked with allele C at another locus. Because of A spreading over the gene pool, C is going to 'catch a free ride' and spread equally, regardless of its phenotypic effect. In this process, genetic variants that are neutral or even deleterious with regards to the fitness of their phenotypic effect can spread through the gene-pool (Ridley, 1993:210).

3.3.6 Conclusion

The mechanism of evolution – selecting random variation in the light of particular environmental conditions – entails that the human mind is not an optimal organ, shaped to produce an accurate and complete representation of reality, but 'merely' an imperfect and particular one, coping with the elements of the environment that mattered for the survival and successful reproduction of our ancestors.

On the one hand, our cognitive abilities are imperfect, because, like any other characteristic, they are the outcome of an undirected, non-teleological process, blindly selecting random variations that prove beneficial, without direction or foresight. Furthermore, they are the result of an evolutionary road that does not undo past acquisitions to start afresh, but merely solves the immediate problems the organism is confronted with, by selecting alterations of the basic building plan, resulting in non-optimal adaptations. Finally, natural selection is not the only evolutionary force. Both the process of genetic drift and genetic hitch-hiking can alter the gene-pool of species in a completely random fashion.

On the other hand, our cognitive abilities are particular, because, firstly, they are shaped in the light of previously acquired traits (path-dependency) and, secondly, they have evolved to cope with a small part of the external world: the elements in our environment that mattered for the survival and reproduction of our ancestors. Vollmer (1975) argues that our intuitive grasp of the world is limited to 'medium dimensions', and evolutionary psychologists point out that our mental wiring – both on an emotional and cognitive level – was shaped to fit the needs of hunter-gathering tribes of the Pleistocene era.

3.4 The lack of external perspective

3.4.1 How does the mind mirror the world?

When Lorenz (1941:124) states that 'just as the hoof of the horses is adapted to the ground of the steppe which it copes with, so our central nervous apparatus for organizing our image of the world is adapted to the real world with which man has to cope' and therefore claims that we can trust our central nervous apparatus to yield an (approximately) accurate representation of the world, he overlooks one crucial element: the lack of external perspective. Indeed, while it is possible for us to see the way other organisms' characteristics are adapted to their environment, we lose this external vantage point from which to appreciate this particular match when it concerns our own ability to mentally represent the world, since we are confined to our mental representation.

We might see, for instance, how the hydrodynamic shape of fish or the front paws of moles, functioning as powerful digging tools, are adapted to their particular environment, but we are in the dark when it comes to assessing how our intuitive representation of the world is adapted to the (real) world we live in. It would, indeed, take another mind to understand how particular characteristics in the world led us, through the process of natural selection, to our particular way of representing it. Furthermore, as pointed out previously, the same environment can lead to a wide variety of adaptations. In this sense, Lorenz's claim that adaptations mirror the environment cannot be taken literally. Do both whale sharks and sea urchins mirror a marine environment? The evolution of our mental abilities might, in this regard, have taken us in a radically different direction, coping in completely different ways with similar environmental conditions.

Therefore, since we have no way of knowing *how* our minds are adapted to the world – how they 'mirror' the world – we cannot make any inference, from the particular view of the world that natural selection has provided us with, as to the states of the external world. We cannot, in other words, derive any knowledge of the Kantian noumenal world (i.e. the world in itself), by means of the phenomenal world (i.e. the world as it appears to us), based solely on the fact that this phenomenal world must be adapted to cope with the noumenal world. The connection between our view of the world and the external world, in this regard, even if we know it must exist, must remain irremediably hidden to us.

Imagine, in this context, that we must infer how a 'real' tree looks based solely on a representation of this tree. Being confined to the representation – let's say a drawing on paper – how could we possibly infer the true qualities of the actual tree it represents? Indeed, what information would we have of its size, its physical properties or even its three-dimensionality?

Furthermore, we would have no means of knowing what elements of the representation are meaningful. In this light, it would be impossible to assess if the visual experience of it matters rather than the tactile experience of feeling the paper, or its smell or taste for that matter. And even if we knew it was a visual representation, we wouldn't have any idea if the empty space around the ink is part of the tree or not, or if the tree corresponds only to the ink lines and not the space in between, and so on. Therefore, being confined to the representation (i.e. the picture of the world created by our perceptual and cognitive apparatuses), we have no way of inferring anything about its object (i.e. the world itself), even if we are guaranteed that there is a correspondence between them.

3.4.2 Analogy with other organisms

Now, assume that some foreign intelligent creature – an 'extraterrestrial scientist' – contemplates our perceptual and cognitive abilities and the view of the world they yield. Benefiting from an external vantage point, this creature could see how our mental representation of the world is adapted to the external world. However, in doing so, could this 'scientist' come to any conclusion other than that we are endowed with a particular way of understanding our surroundings? Just as we conclude, from an external vantage point, that the cognitive abilities of animal species yield a particular grasp of the world – which is certainly adapted to their biological needs and environment, but nevertheless particular – it bears no doubt that this creature too could reach no other conclusion.

Indeed, viewed from an external point of view, our view of the world – our umwelt – would appear just as particular as those of animal species appear to us. Since the cognitive abilities of animal species are every bit as adapted to their ecological niche as ours, the argument that holds that since our cognitive abilities are shaped by natural selection, they must yield an accurate understanding of the world, loses its ground. Adaptations – as pointed out before – do not mirror the world; they 'mirror' very particular aspects of the world in a very particular, contingent way. Nothing suggests that this does not apply to our senses and cognitive faculties responsible for our intuitive representation of the world.

4. Conclusion

Konrad Lorenz (1941, 1973) claims that we can trust our representation of the world to be (at least approximately) accurate, since it is the product of abilities that evolved by natural selection. Indeed, he argues, natural selection cannot do otherwise than to select faculties producing accurate representations over mistaken ones, as the former undoubtedly provide us with a better chance for survival and successful reproduction.

This argument, however, is mistaken. First of all, as Plantinga (1993) points out, natural selection only shapes cognition to yield appropriate behaviour. False ways of forming beliefs that nevertheless lead the organism to act in appropriate ways will obviously not be selected against. In this context, Stich (1990) argues that natural selection will not automatically yield rational thinking. Quite the contrary. Not only will cognitive systems be sub-optimal due to cost/overall fitness trade-offs (cf. Elliot's concept of internal fitness), but even optimally well-designed systems would not automatically be the most rational ones. A less reliable system can indeed provide the organism with a better chance for survival and reproduction due to a discrepancy in costs that different kinds of mistakes entail. This leads to systematic cognitive biases that our minds appear to be subject to, yielding a larger number of mistakes in order to reduce the total cost of mistakes. Munz (1993), on the other hand, points out non-cognitive functions that our belief forming abilities serve. False beliefs, in this light, might prove fitness enhancing in virtue of, for instance, their ability to delimit social groups, providing its members with a clear identity. Irrationality, seen from this perspective, is not failing rationality, but a separate cognitive 'mode' with a purpose of its own.

Moreover, considering the mechanism of evolution, we cannot but conclude that our cognitive abilities are both imperfect and particular. They are, indeed, the result of an undirected, blind process without any foresight that can only select randomly produced variations, on the one hand, and of other random evolutionary forces, namely genetic drift and the process of genetic hitch-hiking, on the other hand. Furthermore, natural selection is path-dependent, building on what has previously been acquired, and adapting organisms to the particular relevant aspects of their particular environment. This leaves them with thoroughly particular characteristics and abilities, as the outcome of both their evolutionary past and the selective forces acting upon them, adapting them to a small subset of elements in the world.

Finally, the Lorenzian argument fails in principle. Indeed, whereas we can see how characteristics and abilities of other organisms are adapted to their ecological niche, we lose this external perspective when it comes to our own cognitive abilities. Not knowing how our mind is adapted to cope with the world, we cannot infer anything about the external world based on our representation of the world. Furthermore, just as we consider other species' perception of the world to be irremediably particular, so – we must admit – would a foreign intelligent creature consider our view of the world as irredeemably particular. Like any other organism's, it is adapted to cope with a particular part of the world in a particular way.

Therefore, through the criticism of the Lorenzian argument provided in this chapter, the conclusion of the previous chapter – namely, that natural selection endowed us with a particular (non necessary and non universal) view of the world – is reinforced. Indeed, from these counterarguments, the conclusion emerges clearly that we cannot infer the truthfulness[16] of our representation of the world from its functionality in terms of survival and reproduction. Furthermore, both the mechanics of evolution and the analogy with other species point towards the constitution of perceptual and cognitive abilities that generate a thoroughly *particular* representation. This, however – as pointed out – does not entail that natural selection only endowed us with false representations, merely that we cannot infer their truthfulness solely on the basis of the fact that natural selection shaped our perceptual and cognitive faculties.

[16] True, accurate or truthful beliefs or representations are, in this context, synonyms. I'm using these terms interchangeably to point out an appropriate relationship between representation and its (external) object, without committing to any theory of truth in particular.

Chapter 3: Transgressing our biological bias

1. Introduction

In chapter 1, I concluded that we are endowed with a particular, species-specific view of the world as the result of our perceptual and cognitive wiring. In chapter 2, I concluded that the fact that these perceptual and cognitive abilities are the product of natural selection does not entail that the view of the world they cause us to have is (even approximately) accurate. This reinforces the conclusion that – in the same way as any other species – we are endowed with a non-universal, non-necessary grasp of the world, as the product of a blind and contingent process. Nevertheless, we are not confined to this particular, species-specific grasp. We can, indeed, 'transgress' the particular commonsense representations yielded by our perception and cognition. We can substitute them with alternative and often contradictory representations that we perceive as more accurate.

In this chapter I enquire about this possibility of going beyond and against our biological bias. How do we do it? What is the source of this epistemic feat? Once this is established, I will look at the scope of these transgressing theories and ask whether, and – if so – to what extent they are limited by our perceptual and cognitive architecture. That will be the subject of the next chapter.

There are three possibilities as to what enables us to transgress our biologically determined view, as opposed to all other animal creatures we are acquainted with. The first possibility is that we are not hard-wired. Nothing is innate in humans and therefore we are not bound to view the world in a particular species-specific way. The second possibility is that those transgressing theories have an innate basis. We have, in other words, much more innate ways of viewing the world than other creatures, which provide us with this distinctive cognitive flexibility. The third and final possibility is that we have similarly biologically constrained views of the world as some other nonhuman animals (as the result of perception and cognitive predispositions), but can perform different cognitive operations on these genetically determined views, enabling us to transgress them.

I will argue that the last possibility is the right one. This, however, presents us with an apparent contradiction. How, indeed, can the mind both be predisposed to view the world in a particular way and at the same time be able to overcome these views? Several theories have been proposed to reconcile innate predisposition and cognitive flexibility. The main proposals of cognitive faculties enabling us to transgress the commonsensical views we hold in virtue of our nature are: mapping across domains, reasoning by analogy, metarepresentational thought, and explicit and conscious representation.

While all these faculties are relevant to the issue at hand, taken individually, however, they fail to account for our ability to transgress our commonsense representations and lead to confusion by covering similar faculties from different perspectives. Therefore, I propose to go about it the other way around. Instead of fitting salient cognitive faculties to human epistemic achievements – and their characteristic feature of transgression – I reason from the ability to transgress in general, and ask what cognitive faculties are required. In doing so, I aim to integrate and complete those proposed cognitive faculties in a framework constituted by the necessary conditions for transgressing innate predispositions, providing an unequivocal and sound account of the distinctive human cognitive ability to go beyond and against its biological bias. However, the questions as to what other faculties underlie these faculties, and how they evolved or how they are physically realised, remain beyond the scope of this chapter.

2. The unique human ability to transgress its biological bias

As discussed in detail in the previous chapters, human beings – in the same way as any other species – are endowed with a particular window on the world, as the result of their nature. We are, in other words, biologically biased to view the world in a particular way. This intuitive view is the result of both our perceptual and cognitive architecture, providing us respectively with an incomplete input dressed in a species-specific phenomenal coat and a predisposition to interpret this input through a framework of innate cognitive 'modules', carving up this phenomenal experience of the world and applying modes of reasoning to it that are notoriously scientifically inaccurate. Natural selection did not, indeed, provide us with an accurate view of the world – as the Lorenzian epistemologists would have it – but, quite to the contrary, with a thoroughly incomplete and particular one, fit for survival and reproduction in particular environments and shaped by a blind process of path-dependent and gradual alterations of previous building plans.

Is it then safe to conclude that we are forever and irremediably trapped in this species-specific, particular representation of the world, with every attempt to represent the world in a more objective fashion doomed to failure? Not exactly. While there is no way around the fact that we perceive and are predisposed to conceptualise the world in this particular way, the rich history of human thought proves that we are somehow able to transgress these commonsense representations. We are, in other words, able to substitute the representations we hold in virtue of our nature, with alternative and often contradictory representations that we consider as more accurate. How else, indeed, could we assert that our intuitive view of the world is flawed?

Homo sapiens, in this regard, is the only species on this planet that goes beyond its 'cognitive niche'[17] or as Uexküll (1909) calls it: its 'umwelt'. We are not bound by the commonsensical, uncritical view of the world that our perceptual and cognitive architecture provides us with. Indeed, we overcome our innate predispositions, come up with different representations and thereby escape this biological determinism that all other creatures are – at least to an important extent – subject to. This ability to go beyond the limits of its given niche is the true wonder of human cognition and the origin of culture, underlying typical human activities and institutions ranging from art to society and from religion to science. Every account of human cognition should therefore account for its characteristic ability to transgress its biologically determined scope on the world.

[17] The term 'cognitive niche', in this context, does not refer to Tooby and Devore's (1987) use of it – i.e. deploying information and inferences rather than particular features of physics in the struggle of life – but to the particular cognitive relation of an organism to its environment as the result of natural selection.

The human mind, in this context, is able to transgress both the input it receives through the senses and its predispositions to conceptualise this input in a particular way. Indeed, we conjecture that the world is not made up of colours and shapes, sounds, odours and tastes, but of atoms which cause us to perceive these qualities in virtue of the particularities of our senses. As pointed out in chapter 1, vision – according to modern science – is caused by photon particles reflecting (or not) off other atoms and hearing is caused by a change in air-pressure. Moreover, both colour and sound are conceptualised as waves: the first as electromagnetic waves and the second as a mechanical wave of air pressure. Furthermore, we infer the existence of unperceivable ranges of potential visual and auditory stimuli. For example, while we can only see light with wavelengths between 375 and 750 nanometer, we have the whole spectrum mapped. Similarly, we are aware of the existence of sound vibrations beyond the audible range of 20 Hz to 20 kHz.

As pointed out, we also transgress our commonsense conceptualisations, explaining entities and events in the world in ways that go against our predispositions. Indeed, contrary to what our commonsense dictates, we conjecture that animate organisms are composed out of purely physical particles and even that mind and consciousness are nothing more than the products of atomic interaction. Moreover, Newtonian physics tells us that moving objects are not moving because of a force acting upon them and are immobile without such an impetus, but that, quite to the contrary, objects will keep going on their trajectory indefinitely unless a force acts upon it. Similarly, Darwinism rejects our essentialist intuition, arguing that all species change over time and are all ultimately descendent from one and the same organism at some point in the history of life. Even more radically, Einstein overthrew our deeply rooted intuition that space and time are two absolute, independent entities in his distinctively counter-intuitive theory of relativity.

Therefore, in order to gauge the scope of human cognition, as I set out to do in this dissertation, it is imperative to analyse this ability to transgress our innately grounded commonsense representations. Indeed, an answer to the question as to what the source of these transgressions is – how, in other words, our mind can reach beyond its intuitive views – will provide us with the necessary tools to tackle the question at the origin of this undertaking: what can we hope to know about the world? While the scope and the eventual limits of these transgressions will be the subject of the next chapter, this chapter is committed to uncovering the cognitive source(s) underlying this distinct human ability to go beyond and against its biological bias.

3. What enables us to 'transgress'?

To what does Homo sapiens owe its distinctive ability to go beyond its cognitive niche? What, in other words, enables us to overcome the view of the world we are predisposed to have in virtue of our nature? There seem to be three possible accounts of our ability to transgress our commonsense representations. The first possibility is that we owe this cognitive flexibility to a lack of innate predispositions. Our minds, under this view, are free of inborn restrictions or predispositions and can therefore produce an unlimited realm of novel and unbounded representations. In contrast with our fellow animal creatures' minds, this hypothesis argues that we are not bound by instincts and come to this world as blank slates on which cultural elements can then be impressed. Under this view, humans – both in the way they represent the world and in the way they generally behave – are all culture, no nature. This hypothesis was put forward by empiricists and, according to Pinker (2002), still prevails in the social sciences. Let's call it the 'blank slate model'.

The second possibility is that our cognitive flexibility is the product, not of a lack of innate predispositions but, on the contrary, of an excess of these predispositions. Flexibility, in this sense, is the result of having a massive amount of different and often contradictory predispositions, yielding unpredictable and diverse representations. Transgressing representations, therefore – i.e. representations departing from commonsense – are themselves the product of innate, content-rich modules (meaning that the content of these representations is determined by those modules). We are, in other words, predisposed to view the world in a very broad realm of different ways, whereas other animal species, being endowed with less modules, view the world through a less 'kaleidoscopic' lens. This hypothesis was proposed by William James (1890), arguing that humans are actually endowed with more rather than less instincts than nonhuman animals and are therefore less conditioned in their behavioural responses. I shall refer to it as the 'Jamesian model'.

The third and final possibility is that our cognitive flexibility is not the result of either a lack or an excess of content-rich modules, but of distinctive cognitive performances. We are, in this light, endowed with a basic set of innate predispositions (as are other nonhuman animals) – resulting in our uncritical, commonsense representations of the world – but are able to overcome these representations, not in virtue of other content-rich modules, but by means of cognitive operations we perform on those representations. Flexibility and creativity in representing the world, from this perspective, is not the result of a lack of restrictions or an excess of starting knowledge, but of how the mind handles its representations. I will call this hypothesis, the 'operational model'.

3.1 The blank slate model

The blank slate model is rooted in the empirical doctrine. Locke (1690:61) famously claimed that 'no proposition can be said to be in the mind, which it never yet knew, which it never yet was conscious of'. The mind, therefore, from an empiricist point of view, is empty at birth and gradually filled with ideas through experience. It does not contain any kind of innate knowledge or domain specific modules, but merely functions as a container for ideas originated in experience. It is, in other words, a sponge, not a computer running different programs.

This empiricist hypothesis of the human mind corresponds with our commonsensical view of knowledge acquisition. Infants are thought to be born with no preconceptions, and during development, their little heads are gradually filled with information they pick up from the environment. They acquire a language and all kinds of beliefs from the people around them and grow up to be members of a particular culture. How they come to see the world, how they behave and even what moral values they adhere to, in this sense, is (almost) purely the product of the culture in which they are born, not their (human) nature. This view is what Pinker (2002) refers to as the 'standard social science model'.

Today, however, this model is largely discarded. It is now commonly accepted, at least in the sciences preoccupied with understanding the working of the human mind, that the mind is not an empty container but contains a number of special purpose programs. This radical turn, known as the cognitive revolution, owes a great deal to Chomsky's work in linguistics and his ground-breaking review of Skinner's (1957) 'Verbal behavior'. While Skinner argues within the paradigm of behaviourism – reducing all mental activity to a set of behavioural dispositions – that language acquisition is the product of conditioning, much as Pavlov's dogs' reaction to the bell, Chomsky (1959:42) counters that children cannot learn language simply through meticulous care on the part of adults, shaping their verbal repertoire through differential reinforcing. Why? Because the linguistic data the child is exposed to can never account for its knowledge of the language. Language mastery, therefore, cannot be the product of having a set of behavioural dispositions – it presupposes the existence of a 'language module', an innate set of grammatical rules and constraints or, as Chomsky calls it, a 'generative grammar', enabling the child to make inferences about the construction of new sentences it has never encountered before. The mind, therefore, cannot be a general purpose processor, but must – at least with respect to language acquisition – be endowed with domain-specific modules.

This new paradigm of human cognition was soon confirmed by empirical evidence from areas such as developmental psychology and cognitive anthropology. Indeed, as pointed out in chapter 1, developmental psychologists Spelke (1991) and Baillargeon (1991) tested 3 to 8 month old infants with regards to their conception of objects and the physical laws governing them and concluded that infants expect objects to be impenetrable by each other, to move along continuous trajectories and to be cohesive. Anthropological research, on the other hand, shows that the human mind is predisposed to think about fauna and flora in a highly structured way. All cultures appear to divide the natural world into a complex taxonomy which incorporates different groups, each further defined in different levels of subgroups (Atran, 1998). Other probable candidates for domains of innately constrained representations are: a sense of numeracy and natural geometry (Spelke 2003), a domain for psychology or theory of mind (Pinker 1997), a domain for facial recognition, and a cheater-detection module (Tooby and Cosmides 1992).[18]

Besides this empirical evidence rejecting the blank slate model, the latter also fails to account for knowledge acquisition in principle. Indeed, as the Latin adage has it: 'ex nihilo, nihil fit'. Nothing can come out of nothing. Therefore, in order for mental representations to take form, we need to presuppose a processing apparatus, or at the very least, an ability to detect similarities. In this regard, Popper (1957) criticises Hume's account of induction. Hume (1739) argues that we believe in law-like regularities in the world based on induction or the repeated experience of similar instances (such as, for instance, the belief that a billiard ball will move when another ball collides with it). Therefore, Hume claims, since there is no guarantee that these experiences will repeat themselves, we cannot be justified in holding these beliefs. While Popper agrees with the conclusion, he rejects the premise. Repetition based on similarity cannot be something 'in vacuo', he claims, it has to be similarity *for us*. Instances can only be similar from a certain point of view. Its recognition presupposes a prior interpretation. Therefore, Popper argues, theory must come before observation. There must, in other words, be a prior framework in which observation can be interpreted, or at the very least, a propensity to expect regularities. An empty mind, in this regard, could never form beliefs, since beliefs can never start with observation alone.

Furthermore, how is a child ever to acquire a linguistic concept such as, for instance, 'cow' or 'house', if it does not have the proper cognitive predisposition to connect the utterance of these words with the objects that are pointed at and not some part of it or some other aspect and, even more importantly, generalise the applicability of these concepts on all similar objects it encounters (the assessment of which already requires a theory, as Popper has shown)? As Quine's (1960) famous 'gavagai' example shows, we can never infer from mere ostention the proper reference with absolute certainty. Indeed, when a tribesman – speaking a tongue we have no knowledge of – exclaims 'gavagai', pointing at a running rabbit, how are we to infer whether he refers to a whole rabbit, some undetached part of it, the colour of its

[18] Whether all of these proposed domains constitute separate and autonomous modules, however, is not my concern. My only aim is to illustrate the generally accepted claim in cognitive and evolutionary psychology that the human mind is endowed with domain-specific knowledge systems, predisposing it to view the world in a particular way.

fur, the temporal presence of rabitness, or so on.[19] The child needs, in other words, a predisposition to zone in on the way natural languages carve up the world in order to acquire concepts. Knowledge acquisition, therefore, presupposes the existence of prior theories, all the way down to a set of cognitive predispositions the newborn child is equipped with, allowing him to couch his experiences in a proper framework.

Finally, the argument that the human mind is endowed with innately constrained ways to view the world becomes obvious when put in an evolutionary perspective. Indeed, since all non-human animals are endowed with such cognitive modules, it is very unlikely that Homo sapiens is the only species which is not. This would mean that in those six million years – not a very extensive stretch of time in evolutionary terms – that separate us from other apes on the evolutionary tree, all these cognitive acquisitions would have been undone. This is extremely improbable. As pointed out before (cf. chapter 2, 3.3.3 Path-dependency), evolution by natural selection does, indeed, not erase previous acquisitions to start from scratch, but can only gradually alter the existing blue-print to fit new requirements. It is, in this regard, beyond any reasonable doubt that our minds still bear the marks of our primate ancestors.

3.2 The Jamesian model

William James (1890) notoriously rejected the traditional view opposing animal behaviour as guided by instincts, and human behaviour as guided by reason. While it is true that humans show more hesitation, reflection and choice in their course of action than animals, he argues that this is, however, not the result of a lack of instincts. Indeed, instincts do not automatically result in invariable motor-reflexes triggered by a particular perceptual stimulus, but can lead to occasional irregularities in behaviour in any animal having a large enough number of separate instincts, which therefore can inhibit the activity of opposite impulses (391). The more instincts, in other words, the more flexible and unpredictable the behaviour of the organism will be. Man, in this light, with its unprecedented variety of behaviour, has – as James put it – 'a far greater variety of impulses than any lower animal' (390).

[19] Quine's example is intended to illustrate the inscrutability of reference or extension - an ontological matter - but can be used for our purposes - the epistemological necessity of innate 'theories' for knowledge acquisition.

A similar position, known as the 'massive modularity of mind hypothesis', has been proposed more recently by the evolutionary psychologists Tooby and Cosmides (1992). They argue that the adaptive problems guiding the evolution of the mind of our ancestors are solved more quickly, efficiently and more reliably by a modular system than by a general purpose system. Indeed, different problems require different solutions, which can be realised much more easily by different modules each coping with one problem, than by a single information-processing design having to cope with all of them (110). Natural selection, therefore, must have pushed for a modular organisation. Our mind, in this light, is best seen as a 'Swiss army knife', consisting of a set of separate, domain specific information processing mechanisms shaped by natural selection to cope with adaptive problems. Flexible human behaviour and cognition, according to Tooby and Cosmides, is the result of a great number of these single purpose modules, not of some general purpose, higher form of cognition.

This, however, cannot be the whole story of the human mind. Indeed, as Fodor (1983) points out, a purely modular mind cannot account for the obvious holistic character of human cognition. This entails that information from the different modules is accessed by the 'central system' of the mind (103). This holistic nature of cognition – this creativity and passion for the analogical – as Fodor puts it (1985:4), is precisely what provides human cognition with its ability to transgress its biological bias. This transgression cannot be accounted for by Tooby and Cosmides' massively modular mind, which would entail that we can only represent the world in ways determined by those very modules. This would leave cognition void of any form of flexibility, let alone creativity.[20]

As Mithen (1996) points out, the modern human mind does precisely what Tooby and Cosmides say that it shouldn't do – it applies reasoning patterns from one module to the subject matter of another. While it is true, Mithen argues, that the human mind is composed of a set of innate, content-rich mental modules – providing us with an intuitive grasp of domains like physics, psychology and natural history – these modules are, however, not isolated. Hunter-gatherers typically reason about the natural world as if it were a social being,

[20] Carruthers (2006), however, defends the massive modularity of mind hypothesis against this critical point, but, in doing so, has to drop the characteristic encapsulation of the modules in his model of mind. In doing so, he backhandedly confirms that cognitive flexibility entails the flow of information across the boundaries of the different modules.

conceiving of the forest as a parent and animals as kin, hereby imbuing the world with human qualities such as will and purpose (47). Similarly, scientists view biological organisms as mere physical objects, composed out of atoms, rather than distinguishing sharply between animated and inanimate substances, applying the reasoning pattern of the module of biology to the former and that of physics to the latter. In short, although our mind is composed of different 'blades' shaped by natural selection, we do not necessarily apply the blades to their fitting problem, as Tooby and Cosmides would have it.

The argument that our cognitive flexibility and creativity is not the product of our innate endowment is also supported by empirical evidence from a different angle. Spelke (2003) compares the innate modules, or core knowledge systems, as she refers to them, of humans and nonhuman animals. The first is our core system for representing material objects. As pointed out, infants represent objects as cohesive bodies, moving along continuous trajectories and interacting with each other only by contact. This intuitive grasp of objects, however, is also found in a variety of nonhuman animals such as monkeys and even newly hatched chicks (282). Similarly, our sense of natural geometry, enabling us to navigate through spatial layout, and our sense of numeracy, enabling us to represent small numbers and estimate approximal numerical magnitudes, is found both in infants and in other primates (297). The considerable cognitive gulf between human and nonhuman animals, therefore, cannot be the product of different sets of starting knowledge. The distinctively human ability to go beyond its biological bias, in this regard, cannot be the product of its innate endowment.

Finally, an evolutionary argument can be brought to bear here. As Tomasello (1999) points out, the 6 million years that separate human beings from their primate cousins is a very short time in evolutionary terms – not enough, Tomasello argues, for natural selection to have produced a whole different suite of cognitive modules one by one. Moreover, he continues, paleoanthropological research suggests that for all but the last 2 million years the human lineage showed no signs of anything other than typical great ape cognitive skills, and the first dramatic signs of species-unique skills only entered the scene in the last quarter of a million years with modern Homo sapiens (4). This reinforces the argument that the source of the cognitive divide between human and nonhuman does not reside in a radical difference in their innate, content-rich cognitive endowment. This brings us to the final hypothesis.

3.3 The operational model

Since our cognitive ability to transgress our commonsense representation is not the product of our innate endowment – since, in other words, we have no more information to start with than some nonhuman animals who do not transgress their innate representations – we have to account for this ability in terms of operational cognitive faculties rather than content-rich modules (either by postulating the absence of restrictions or the excessive presence of those modules). This is what the operational model proposes. According to this model, our biological bias is the product of our innate endowment, i.e. our senses and cognitive wiring. This innate endowment saddles us with particular intuitive theories. Nevertheless, we are the only species on this planet which is not bound by these representations. We can alter them, developing new ways of viewing the world and escaping our cognitive niche as the result of cognitive operations the mind performs on its representations.

The hypothesis set forward by the operational model allows for the co-existence of the necessity of epistemic constraints for learning (cf. Popper's argument that theory must precede observation and that the evolution of knowledge must eventually by grounded in innate dispositions) and the characteristic flexibility and holism of human knowledge (cf. Fodor's argument against the massive modularity of mind hypothesis). Furthermore, it agrees with both the empirical evidence from developmental psychology and anthropology, uncovering the existence of innately determined representations of the world, and the obvious transgression of these representations in the history of human knowledge. Finally, it is on par with evolutionary considerations, allowing our cognitive architecture to still bear the marks of our primate ancestors, and not requiring the independent evolution of a whole new set of content-rich modules in the short period of time (in evolutionary terms, that is) that separates the human lineage from other apes. Human cognition, therefore, is characterised both by a starting set of innate predispositions, and an ability to transgress the representations shaped by these predispositions.

How is this possible? How can the mind both be predetermined to yield a particular set of representations, and at the same time able to overcome those restrictions and conceptualise the world in (radically) different ways? How, in other words, can we transgress the representations we are bound to have in virtue of our biological nature? This question has led to a wide variety of theories, reasoning from a nativist stance – claiming, as I do, that the human mind possesses a variety of special purpose, domain-specific knowledge systems – and proposing cognitive faculties that enable human beings to conceptualise the world in ways that go beyond and against their innate predispositions. These theories about 'what makes us smart' take up the challenge to reconcile the premise that human cognition is based on innate knowledge systems or modules and the obvious flexibility displayed by the mind in its conceptualisations. Let's have a look at the main proposals.

4. Theories about what makes us smart

4.1 Mapping across domains

Carey and Spelke (1994) attempt to explain conceptual change in cognitive domains, as evident from the history of science, within a framework of domain-specific cognition. Human minds, they argue, are endowed with innate systems of knowledge, each pertaining to a particular domain – such as, for example, other minds, physical objects or numbers. According to these authors, such domains comprise a distinct set of entities and phenomena (e.g. the innate knowledge system of physics applies to material bodies and their behaviour) and a number of core principles (in the case of physical objects: cohesion, continuity and contact). Learning, from this perspective, consists of an enrichment of those core principles through experience. Therefore, a high degree of universality in representations is to be expected in the domains for which humans possess innate core knowledge systems. Nevertheless, scientific theories demonstrate that conceptual change in those domains is both possible and actual (Carey and Spelke, 1994:169). Core principles are overridden and new principles adopted, leading to widely diverging theories.

According to Carey and Spelke (1994), conceptual change in those innate domains is the result of 'mapping across domains' (180). This happens when the core principles of one system are applied to the set of entities of another system, thereby escaping the principles that naturally – i.e. in virtue of our nature – fit these entities. By devising and using systems of measurement in physics, for example, scientists create ‘a mapping’ between the core knowledge system of numeracy and that of physics. Therefore, the principles governing the behaviour of physical bodies are no longer those of cohesion, continuity and contact but the core principles of the system of numeracy, such as 1-to-1 correspondence, succession and the like.

Mithen (1996) takes on a similar viewpoint from a different perspective. Reasoning from an archaeological background, he attempts to explain the cultural explosion in the transition between the upper and middle Palaeolithic eras (approximately between 60 and 30 000 years ago). According to Mithen, this ‘big bang of human culture’, featuring the birth of art, religion and complex hunting strategies, is the product of a final, major re-design of the mind in human evolution (153). Whereas our human ancestors evolved encapsulated, domain-specific knowledge systems in areas such as social intelligence, natural history or biology, and technical intelligence, these specialised intelligences were now no longer working in isolation. The resulting cognitive fluidity or mapping across domains, characteristic of the modern human mind, provides – according to Mithen – the foundation of scientific endeavours and the distinctive human ability to transgress the contents of its innate, domain-specific knowledge systems.

The view that the human mind integrates the content of different domains, underlying Carey and Spelke’s (1994) and Mithen’s (1996) account, is a recurring theme in the cognitive sciences. According to Fodor (1983), the mind consists of fast, mandatory, encapsulated and domain-specific input systems or modules, on the one hand and a central system which is slow, non-mandatory, non-encapsulated and domain-general, on the other. This system integrates the outputs of the modules and provides human cognition with its characteristic holism and creativity. As to the nature of this system, Fodor remains mute, considering it an irresolvable mystery. Carruthers (2006), on the other hand, postulates a massively modular mind. He accounts for its flexibility and creativity by distinguishing between two reasoning systems. The first corresponds to the processing of the modules: it is arranged in parallel and operates swiftly and unconsciously. The second supervenes on the activity of those systems: it is realised by mental rehearsal in general and inner speech in particular and operates more slowly and consciously. It integrates, in other words, the content outputs of the various modules, overriding the results of the first system (254). Boeckx (2010:128), finally, compares the cognitive ability to cross the boundaries of modules – underlying our rudimentary theories of the world around us – to mixing apples and oranges to form what he calls 'delightful cognitive cocktails'. Like Carruthers, he points to natural language as a content integrator of the different outputs of the modules.

4.2 Analogy and metaphor

Gentner (2003) argues that higher-order cognition is the product of our capacity for analogy. In this regard, our ability to draw abstractions from particulars, to maintain hierarchies of abstraction, to reason outside the current context, to compare and contrast representations and to project further inferences, among other distinctively human abilities, is seen as the product of our ability to learn by analogy. This inborn faculty is, according to Gentner, multiplied by the possession of relational language, which both invites the learning of relational concepts, and provides cognitive stability once they are acquired (195-196). Therefore, while we possess – in the same way as other animal species – a basic set of cognitive constraints in the form of attentional biases and learning propensities, we also have the possibility to go beyond these biases by means of what Gentner calls 'structure-sensitive comparison processes'. In other words, the capacity to detect similarities in abstract relational structures enables us to make inferences that transgress our innate set of starting knowledge (227-228).

This process of abstraction by analogy becomes evident when considered from a different angle: the pervasiveness of metaphor in language. As Lakoff and Johnson (1980) point out, metaphor is not just a device of poetic and rhetorical flourish, nor even a purely linguistic matter, but is at the core of our everyday thought and action (3). Language, in this context, is merely a body of evidence for the way we think, showing us that human thought processes are largely metaphorical. The concept of 'argument', for example, is structured, according to Lakoff and Johnson, by the metaphor 'argument is war'. We call some claims indefensible, we attack weak points of an argument, we can be right on target, demolish an argument and, of course, win or lose it. Similarly, 'time is money': we waste, spend and save time, chores cost time, we invest time in each other, and we run out of time (5-8).

Furthermore, on a more fundamental level, spatial orientation metaphors organise entire systems of concepts with respect to one another. More is up, less is down (e.g. numbers go up and down); happy is up, sad is down; good is up, bad is down (e.g. it's going downhill or it's going up, he's rising to the top or tumbling to the bottom); virtue is up, depravity is down (e.g. having high or low standards), and so on. The same goes for front-back, on-off, in-out, centre-periphery, and near-far. Similarly, the metaphor of physical objects and substances structures a whole realm of concepts. As Lakoff and Johnson (1980:25) point out, this allows us to pick out parts of experiences and treat them as discrete entities or substances. Events, activities, emotions and ideas are conceptualised as entities, allowing us to quantify them (e.g. I feel too much anger), identify aspects (e.g. the ugly side of his personality), identify causes (e.g. the pressure of his responsibility caused him to resign) and set goals (e.g. he went off in search of fame and fortune). Moreover, even our most basic concepts are often metaphorically structured. Our concept of time, for example, is structured by the metaphor of space – e.g.: he arrived at 1:30, worked through the night, looks forward to tomorrow, leaves the past behind, and so on (St Clair 2007, Pinker 2007).[21]

Pinker (2007:233) argues that the concepts of substance, space and time (rooted in the metaphor of space) and causation (rooted in the metaphor of force) are 'the substrate of our conscious experience'. They are, in other words, the building blocks of our reasoning, giving us the tools to conceptualise about the physical and social world, and, most importantly, they are the source of the metaphors by which other spheres of life are comprehended. Metaphors in language, in this regard, are proof of the way our mind co-opts reasoning patterns that are grounded in our innate knowledge systems to reason about other, abstract domains.

Johnson (2007), in this context, grounds meaning and the nature of abstract thought in image schemas, arising in our perception and bodily movement. Abstract concepts and thought do not constitute a wholly different kind of logic, but are an extension of spatial-bodily concepts. Thought is not disembodied, Johnson argues, with Lakoff (1980), but structured by our sensorimotor schemas and extended by means of conceptual metaphor (Johnson, 2007:180-181). Precisely this ability, grounded in our innate cognitive architecture or bodily generated

[21] This, however, does not imply that space and time cannot be represented independently in the brain, but merely that they are structured by the brain in a similar fashion (cf. Kemmerer, 2005).

schemas – as Lakoff and Johnson (1980, 2007) argue – to extend reasoning patterns by conceptual metaphor to abstract domains for which we have no knowledge-system, enables us to grasp the latter in terms of the former. This process, as pointed out, leaves metaphors as a tangible trace in language. For, as Pinker (2007) puts it, language is the mirror of thought.

4.3 Explicit and conscious representations

Karmiloff-Smith (1992), studying the human mind from a developmental perspective, argues that we can go beyond domain-specific constraints by a process called 'representational redescription'. In this process, 'information already present in the form of implicit information *in* the mind, becomes explicit knowledge *to* the mind' (18, her italics). While implicit information is embedded in procedures, isolated from other parts of the cognitive system, and merely enabling us to respond to the environment, explicit representations are available to consciousness.

In this light, the innate, intuitive grasp of objects and the physical laws governing them - which Spelke (1991) and Baillargeon (1991) discovered in infants – are present, according to Karmiloff-Smith, in the form of procedures triggering responses to environmental stimuli. They are implicit representations. When this embedded information becomes accessible through the process of representational redescription, children become – in the words of Karmiloff-Smith – 'little theorists' (78). The core principles (i.e. cohesion, continuity and contact) are now represented explicitly, encoded in a format usable outside normal input/output relations and available to verbal explanation. Karmiloff-Smith argues that this does not happen exclusively through conventional learning, by acquiring the representations in linguistic form from parents and educators. While some of the representations might be acquired this way, other theory building occurs by this internal process of representational redescription. The human mind is, in other words, endowed with a mechanism that can bring implicit representations to consciousness, taking them as objects of cognitive attention and therefore enabling it to manipulate them. This, Karmiloff-Smith concludes, permits the mind to extend well beyond its environment and underlies its distinctive creativity (192-193).

Similarly, in recent work, Carey (2009) invokes the mechanism of 'Quinean bootstrapping' to account for the conceptualisation of genuinely new representations. According to Carey, new and 'richer' representations can arise out of representations with more limited expressive power. Invoking the case of a child's acquisition of the concepts of natural numeracy – which exceeds the content of the core knowledge-systems at our disposition – she argues that memorisation of the counting sequence by rote eventually enables children to correlate these number words with matching number sets, making them 'cardinal principle knowers'. In this regard, the acquisition of number words, which explicitly represent numerical values, provides the child with a richer representational medium than the core knowledge systems it started out with.

According to Karmiloff-Smith (1992), explicit representations are intrinsically linked to conscious accessibility. Schacter (1989) draws a distinction between implicit and explicit along this line, with regards to memory, and elaborates upon it. Memories, he argues, are implicit when they facilitate performance of a particular task that does not require conscious recollection of previous experiences, and explicit when performance does require conscious recollection. Consciousness, in this context, refers to 'a person's ongoing awareness of specific mental activity' (356). Therefore, explicit representations are representations of which the beholder is *aware*. In virtue of this awareness, he or she explicitly represents this representation. This brings us to the final proposal of humankind's ability to go beyond its commonsense representations.

4.4 Metarepresentational thought

A metarepresentation is a representation of a representation. We can, however, distinguish between two different kinds of metarepresentation. The first comes from psychology and refers to the possession of a theory of mind. From this perspective, a person holds a metarepresentation by representing another person's representation. For example, Mary holds a metarepresentation when she sees Tom looking for his coat in the closet and infers that he believes his coat is in the closet. In doing so, she forms a representation of Tom's belief or representation, i.e., that the coat is in the closet. The second kind of metarepresentation refers to a representation of *one's own* representations.[22] This is the kind of metarepresentation relevant to the question as to 'what makes us smart'. According to Stanovich (2004), the possession of the ability to represent one's own representations is what separates human from nonhuman animals. It gives rise to the self-critical stances that are a unique aspect of human cognition (1264). It enables us, in other words, to form beliefs about our own beliefs. How reliable are they; upon what are they grounded, etc?

Sperber (2000) agrees with Stanovich that the ability to form metarepresentations is one of the distinctive human cognitive abilities. Like echolocation for bats, he argues, the capacity to metarepresent is both unique to humans and crucial in explaining their behaviour (117). While animals may have some rudimentary capacity to metarepresent in the psychological use of the term (e.g. detecting that a conspecific wants to mate or fight), these forms of metarepresentation, according to Sperber, lack both compositionality and recursion. They can only metarepresent a short and fixed list of representations. Humans, on the other hand, can metarepresent an unlimited amount of representations. According to Sperber, this requires a whole new level of cognition. The mental ability to represent does, indeed, not imply the ability to represent those representations. They would remain hidden to the beholder, unless there is something that renders them tractable and therefore cognitively accessible (118-121).

According to Dennett (2000), this something is provided by the encoding of representations in language or other tangible media of representation (e.g. drawings, writings, etc.). The obvious route to true, genuine metarepresentation – in this case, the self-conscious representing of one's own representations, which Dennett calls 'thinking about thinking', not to be confused with the representation of another person's beliefs (cf. the first kind of metarepresentations) – is, he argues, from the outside in. It begins with overt use of public symbols – i.e. the acquisition of natural language – and creates practices that can later be internalised, providing us with the necessary tools to think about thinking (21). Indeed, our possession of a medium in which representations can be couched (such as that provided by natural language) enables us to form what Dennett calls 'florid representations' as opposed to 'pastel representations'.

[22] It is very plausible that the ability to metarepresent in this second sense evolved from the previously acquired ability to metarepresent in the psychological sense – i.e. the possession of a theory of mind. Selective pressure on social intelligence is, as Mithen (1996) points out, suspected to be the motor behind the evolution of human intelligence. A discussion on the origin of our ability to metarepresent, however, is beyond the scope of this chapter.

While the latter are merely unconscious guides to behaviour, the former are deliberate, knowing and even self-conscious ways of representing. In the terminology of Karmiloff-Smith (1992), they are explicit representations. These florid representations then, according to Dennett (2000), are truly metarepresentational in kind, prompting him to exclaim that there can be 'no florid representation without metarepresentation' (19).

5. Necessary conditions for transgressing

While all of the proposed accounts of what distinguishes human cognition from other forms of cognition, with its characteristic flexibility and non-encapsulation – underlying its ability to transgress innately grounded core intuitions about the world – indicate important aspects of human cognition, I argue that taken individually they cannot account for it. Furthermore, they lead to confusion by highlighting similar faculties from different angles and depicting them via different terminology. In order to bring this incompleteness to light and dispose of the terminological confusion, I propose to reverse the sequence of reasoning. Rather than analysing human cognition and fitting salient faculties to our ability to produce theories which depart from our innate predispositions, I will take this ability to transgress our biologically based views as the starting point and analyse what faculties are required to achieve this.

My aim, in other words, is to uncover the necessary and jointly sufficient conditions for transgression. Based on this analysis, I will then point out that the proposed cognitive faculties of 'what makes us smart' are not sufficient to account for this ability to transgress. My analysis, in this regard, takes a non-empirical approach, since its goal is to elucidate what is necessary for transgression in general, not merely how humans achieve this cognitive feat. To put it in another way, I consider the task at hand and ask what is logically presupposed to achieve this. This 'external' analysis, uncovering the necessary and jointly sufficient conditions to form transgressing representations, will then yield a framework in which the current (empirical) accounts of cognitive faculties that 'make us smart' can be integrated and completed where necessary. This, I hope, will provide us with a more complete and unequivocal account of the distinctively human ability to transgress the core intuitions it holds in virtue of its cognitive nature.

For exposition's sake, I will illustrate my argument with a thought experiment. My use of a thought experiment, in this regard, is not necessary, but serves an illustrative purpose. Its goal is to provide the reader with a vivid and clear illustration of what it takes to transgress biologically based views of the world. I could, in other words, bypass the story-telling in my analysis. In my opinion, however, the following thought experiment both enhances the clarity of an otherwise rather abstract exposition and enables the reader to engage intuitively with the matter at hand.

5.1 E.T.s on an icy planet

Imagine extraterrestrial organisms living on a planet at some constant distance from a star, their sun. The planet revolves in such a way around its axis and around the star that the same side is always exposed to the light, while the other is always couched in darkness. Our extraterrestrials live on the side exposed to the sunlight and heat and, therefore, in constant daylight. Furthermore, there are no climatic changes whatsoever: the temperature remains constant at 5° Celsius and the sun is never obscured by clouds. The landscape of this planet is filled with huge ice caps. Because the temperature is always above melting point, those ice caps are slowly melting. How the ice got there in the first place can, of course, for the purposes of this thought-experiment, be ignored.

These extraterrestrial organisms – let's call them E.T.s – are endowed with a sense of vision. Furthermore, they possess a concept of time and causality similar to ours. Based on this input (their visual apprehension of ice becoming water) and their cognitive architecture (their predisposition to situate this event in time and look for a causal explanation), they are predisposed to think that ice has an inherent quality of becoming water over a certain amount of time. This representation or belief is, in other words, part of their folk physics. Not once, given the climatic conditions of this planet, has this expectation been violated. However, over time a bright E.T. comes up with an alternative explanation: it is not the inherent nature of ice that causes it to become water, but it is the sun which causes ice to become water. My question now is: what cognitive faculties does it take to allow for this radical transgression of the commonsense representation of the E.T.s? How, in other words, can our bright E.T. come up with this 'Copernican revolution', given its input (vision) and conceptualisation faculties (a notion of time in which melting takes place, and a notion of causality)?

5.2 What cognitive faculties would E.T. need to transgress its commonsense view?

First of all, since transgressing its commonsense view not only implies that this view is substituted by an alternative view, but that the latter is perceived as 'better', or epistemically more desirable than the former, our E.T. would have to represent both its commonsense view – i.e. that ice becomes water because of its inherent nature – and its transgressing view – i.e. that ice becomes water because of the sun – in order to compare both. If this representation is implicit – i.e. not represented itself, but merely underlying its expectations – E.T. will never be able to assert that one representation is better than the other, nor even be aware of the two distinct representations, for that matter.

Furthermore, E.T. would need an epistemic goal: a disposition to look for truth or for an accurate description of its external environment; in this case, the transformation of ice into water. Without this epistemic goal, transgression is not possible. Indeed, remember that transgression entails not merely a shift in belief – which could occur without the cognitive creature being aware of it and without an epistemic orientation – but a perception of the transgressing representation as 'better' than the commonsense view, entailing a comparison of both representations in virtue of an epistemic goal.

Thirdly, short of divine inspiration, E.T. must come up with a different representation based on the input it receives and the conceptual tools it possesses. As I have outlined, it possesses a visual input: it sees the ice, the water and the sun; and it has a conceptual architecture representing these entities in a framework of time and causality. This has led it to believe that ice becomes water *because* of its own nature over a certain amount of *time*. In order to produce the alternative representation that properties of the sun, instead of the ice itself, cause the ice to become water over a certain period of time, it has to combine the representations drawn from its visual and conceptual resources in a different way. In this case, the causal connection between the perceptual input of ice and water has to be attributed to a foreign element – the sun – instead of ice, the object of transformation itself.

This ability to recombine elements from input and conceptualisation into a new representation further requires that E.T. not only represents its commonsense representation and its transgressing representation as such, but represents the parts of this representation as well. It has to hold the representation of 'ice', 'water' and 'causation in virtue of' as separate conceptual building blocks in its mind. This allows E.T. to divorce its representation of 'causing to become water', from its representation of 'the nature of ice', on the one hand. and to form a new representation by reassembling elements from its input (ice, water, sun) and elements from its conceptual resources (in this case, the causal connection), on the other hand. Indeed, since it represented the parts of its commonsense representation, driving the proverbial wedge between them, these parts can now be reattached with another element it represents (i.e. the sun).

Finally, E.T. will need a way to assess whether its new representation (i.e. ice becomes water because of properties of the sun) is preferable in terms of its epistemic goal (i.e. truth or accuracy of description) than its previous commonsense representation (i.e. ice becomes water because of properties of the ice itself). To make this assessment, E.T. needs two things: data demarcating both representations, and epistemic criteria in virtue of which one representation accounts for the demarcating data in a better way than the other (in terms of the epistemic goal, that is).

In this case, our bright E.T. could have noticed that when casting shade over the ice, the ice stopped melting (remember the constant temperature on the planet in the exposure of the sun is 5° C; when the sunrays are blocked, however, the temperature tumbles to – 5° C).[23] This leaves the choice of sticking to the old representation while accommodating the new finding (i.e. ice becomes water because of properties of the ice itself; however, shaded ice does not become water) or accommodating these findings in the new representation (i.e. ice becomes water because of some properties of the sun; therefore, when the sunlight is blocked, so are

[23] Let's assume that there isn't anything casting a constant shadow on the ice caps, and therefore that the effect of shade on ice was never revealed before.

those properties of the sun, which causes their causal effect to vanish).[24] How can E.T. assess wheher the latter representation fits its epistemic goal of accurate description better?

To make this assessment, E.T. would need an epistemic orientation. Such an orientation is provided by a set of values or criteria which can rank different representations on a scale determined by its epistemic goal (i.e. a scale ranking from less to more accuracy of description). Those values, in this case, could be explanatory scope (the latter representation explains more since it offers a causal reason why non-shaded ice does not become water, whereas the former does not); simplicity (rather than attributing causal powers to one state of ice – i.e. non-shaded ice – and not to another – i.e. shaded ice – E.T. can attribute an unchanging causal power to the sun); and coherence (while the commonsense representation entails that the nature of ice both causes and does not cause it to melt depending on it being shaded or not, the transgressing representation does not harbour such a contradiction).[25]

5.3 Framework of necessary cognitive faculties for transgressing

There seem to be three major cognitive faculties at play which enable E.T. to come up with a transgressing representation of what turns ice into water. The first is the ability to represent representations and their parts. Indeed, in order to compare both representations, E.T. needs to represent both. Furthermore, it has to represent the different parts of this representation separately. In this case, these parts are: 'ice becomes water', 'because of', and 'qualities belonging to ice'. This, as pointed out, is necessary for E.T. to recombine these parts in a different fashion, allowing it to conceptualise a different explanation of the subject matter (cf. the second condition).

A second necessary condition for E.T. to come up with a representation transgressing its commonsense representation is the possession of a way to recombine the information it

[24] Typically, these data will be at the start of the exploration leading to a new representation, fuelling doubt and directing the conceptual recombination of the available building blocks. This account, however, is not a chronological account of how new representations are typically formed. Its sole aim is to provide us with a clear overview of what cognitive operations are needed to transgress a representation anchored in innate pre-dispositions.

[25] Those values are human epistemic values. Transgression, however, by no means entails the application of these particular values or criteria, but merely the possession of an epistemic orientation – i.e. an epistemic goal and criteria for realising this goal.

gathers, through its visual input and the conceptual resources at its disposal, into a new representation. Indeed, as pointed out, E.T.'s new representation has to be composed out of elements it gathers from its input and the conceptual tools it possesses. Since representations don't materialise out of thin air, any new representation has to be accounted for in terms of the information that can be drawn from the world, and the ways in which this information can be processed.

A third and final necessary condition is the possession of an epistemic goal and epistemic values or criteria. Without an epistemic goal – a view of what makes a representation desirable, such as truth or accuracy of description – one cannot propose alternative representations which are perceived as epistemically more desirable. This goal or ideal is, indeed, necessary to provide an axis on which different representations can be compared – necessary, but not sufficient. For this comparison to take place, there need to be epistemic values or criteria determining the relative proximity of the two representations with regards to this ideal. A disposition to look for true representations is, indeed, vacuous without some criteria that make a representation more or less truthful.

In the case of our bright E.T., those values were explanatory value, simplicity and coherence. As pointed out above, it could only apply those after integrating demarcating data into the equation. This, however, since it was gathered by its visual input, can be accounted for by the previous faculty enabling it to represent information gathered from its input, and to recombine it with representations drawn from its conceptual framework. Moreover, one can imagine a case in which there would be no need for empirical data demarcating both representations; one representation being preferable merely in virtue of epistemic values. For instance, when two representations account for the same phenomenon, but one does so in a more parsimonious way, this representation can be seen as better, yielding to Occam's razor and satisfying the value of simplicity. The better representation, in this case, would typically contain fewer elements and not require additional data but, on the contrary, would erase data from the equation.

Without this last, crucial condition – i.e. the possession of an epistemic value system – alternative representations would merely be random proposals, with nothing to determine whether one representation is more desirable than another. Indeed, the very concept of transgression presupposes the existence of such a system. A representation can only be perceived as transgressing commonsense when it is perceived as more desirable than its commonsense counterpart. For our E.T.s, just as for us, this increased desirability comes from a sense of increased accuracy or truthfulness. In this context, except for extreme cases of scientific relativism, it is commonly accepted that scientific accounts present us with better descriptions of the world than uncritical commonsense assumptions. They are either considered as more truthful (scientific realism), or at least as more useful (instrumentalism).

6. Integrating human cognitive faculties in the framework

6.1 Representing the representation and its parts

As pointed out, in order to compare two representations, we need to represent those representations. If our representations are merely implicit, underlying our behaviour without us representing the representations themselves, how could we assert that one representation is more accurate than the other? Moreover, as pointed out, if representations cannot be carved up into different parts, each represented individually, we have no means of representing the subject matter differently by recombining the different parts.

Karmiloff-Smith (1992) argues that we have the ability to bring implicit representations to consciousness, representing them explicitly. Explicit, in this sense, means available to conscious awareness.[26] The question remains, however, as to what exactly is available. Dienes and Perner (1999) distinguish different kinds – or as I would put it: levels – of explicitness. The content alone of a representation can be represented explicitly (the cat is on the mat); both the content and the attitude can be represented (knowing or believing that the cat is on the mat); and finally the content, the attitude and the holder of the representation can be represented (it is me who believes that the cat is on the mat) (737). Nevertheless, they argue, under a common understanding of the term 'conscious', a representation counts as

[26] At least, this is the case in its ultimate stage, because Karmiloff-Smith conjectures about several stages of gradually increasing explicitness.

conscious only when its content, the attitude and the holder or self can be represented consciously (740). How, indeed, can one be said to be aware of having the representation that the cat is on the mat, without being aware that one knows, sees or believes this, and that, quite obviously, it is oneself who holds this representation? An explicit representation, therefore – in the sense of a consciously available representation – is a representation of oneself representing a content in virtue of a certain attitude. It is thinking about thinking, as Dennett (2000:21) puts it; it is thoroughly metarepresentational in kind.[27]

In this sense, Karmiloff-Smith's (1992) explicit representations and Stanovich's (2004), Sperber's (2000) and Dennett's (2000) metarepresentations refer to one and the same ability of representing one's own representations.[28] They are representations of a higher order because their object is a representation and not something external to the mind. This higher-order cognition, representing the representations we hold in virtue of our biological nature (i.e. our senses and our cognitive architecture), is a first major step towards transgressing the commonsense beliefs we hold. As Stanovich (2004) points out, this ability is distinctively human. While other animals might be able to metarepresent to a certain degree in the psychological use of the term (cf. 3.2 Metarepresentational thought), they cannot be said to represent their own representations. They do not, in Dennett's (2000) terms, think about thinking.

Moreover, when representing one's own representations, the parts of those representations are necessarily represented as well, for how are we going to represent the visual representation 'the cat is on the mat', without being able to represent 'cat', 'mat' and 'on' as separate aspects of this representation? Indeed, while, for example, the implicit visual representation 'the clouds are dark' will simply trigger the instinctive reaction of an organism to seek shelter from the coming rain, its explicit counterpart 'the clouds are dark, therefore it will rain', obviously requires us to represent all elements gathered from input: clouds, dark and the inference from conceptualisation: 'therefore it will rain', individually.

---

[27] It is metarepresentational in the sense of representing one's own beliefs, not representing someone else's beliefs or possessing a theory of mind (cf. the different kinds outlined above – 3.4 Metarepresentational thought).

[28] This does not, however, entail that they agree on which faculties underlie the ability to form metarepresentations, merely that all these accounts point to the ability to represent one's own representations. The question as to which faculty underlies the human ability to form metarepresentations, or how this faculty evolved, is beyond the scope of this chapter.

While Dennett (2000) argues that these 'florid representations' are couched in natural language, Karmiloff-Smith (1992:15) points towards the redescription of implicit representations in different representational formats that are ultimately available to verbal report. Explicit representations, on both accounts, therefore, can in principle be encoded in natural language,[29] a system which constructs representations based on the combination of explicit units of meaning. This provides an explanation as to why, when we represent explicitly, we necessarily do so by representing the parts explicitly. Indeed, either natural language itself functions as the system in which representations can be couched, and therefore made explicit, as Dennett suggests, or these representations are encoded in a format close enough to natural languages to be verbalisable, which entails that this format shares its basic structure with language, as Karmiloff-Smith suggests. This aspect of explicit representation underlies the possibility of reassembling those separately tagged elements in different configurations, as I will discuss the next section.

6.2 Variation through recombination

Merely representing representations does not by itself enable us to transgress these representations. This requires a cognitive faculty that can produce variation. Since we cannot access different ways of drawing input from the world or different ways of conceptualising this input, we can only come up with alternative representations by recombining elements we gather from input and from our conceptual modules in different ways, to form – as Boeckx (2010:128) calls them – 'delightful cognitive cocktails'. We can both explain a familiar subject matter in terms of the principles governing conceptual domains other than the one that we are predisposed to apply to the subject matter, and direct our mind to new subject matters (for which we have no predisposed conceptual grasp), understanding them in terms of familiar domains.

This ability, referred to as 'mapping across domains' or 'reasoning by analogy', enables us to apply our sense of numeracy to the domain of space and time, representing delimited parts of

[29] Karmiloff-Smith (1992:22-23), however, allows for levels of explicitness of representations which are not yet available to verbalisability. She does not, therefore, reduce consciousness to verbal reportability. However, at the ultimate level of explicitness, the representation is encoded into a format which, she hypothesises, is close enough to natural language for easy translation into communicable form.

space and stretches of time with numbers, or to apply our module for reasoning about animate creatures to inanimate objects, investing rocks, trees and the like with spiritual essences, as is the case in animistic religions. Furthermore, it enables us, or rather Harvey in this case, to conceive of the heart as some sort of mechanical, pump-like device, or Bohr to view the structure of an atom as that of a solar system (De Cruz and De Smedt 2007), or, even more fundamentally, it enables us to view time in terms of space, or causation in terms of force, as St. Clair (2007) and Pinker (2007) point out. It allows us, in other words, to think differently about subject matters and to think about different subject matters than the ones we are predisposed to think about.

Language, I have argued, following Lakoff, Johnson (1980) and Pinker (2007), provides us with tangible proof of these cognitive operations. This ability to produce variation by recombination becomes evident when we look at language's compositional character. By recombining words to form new sentences, there are indeed no limits to the amount of sentences with distinct meaning that we can create. This implies that there are no limits to the amount of representations the human mind can come up with. We can, in other words, endlessly recombine those building blocks we gather from perception and conceptualisation into new representations. Moreover, our ability to extend representations to previously unknown domains is evident when we look at the metaphorical character of language. As shown by Lakoff and Johnson (1980, Johnson, 2007), our language is pervaded by metaphors, pointing towards our ability to map body-based, sensory-motor source domains – i.e. innate reasoning patterns – onto new abstract target domains, by means of conceptual metaphor (Johnson, 2007:177).

6.3 Epistemic value system

As argued above, without an epistemic value system – i.e. an epistemic goal and values or criteria for realising this goal – there can be no transgression of biologically based views, since this entails not only the substitution of commonsense views with alternative ones, but also the perception of the latter as epistemically better than the former. Therefore, the human mind has to possess an epistemic goal and epistemic values in order to rise above its uncritical assumptions.

Such a goal, it bears no doubt, is our predisposition to look for truth. Papineau (2000) argues that the search for truth is an innate drive, much like hunger and the desire for sex. It is, in other words, part of our innate endowment: a product of natural selection that increases our chances of success in our practical projects and thereby boosts our biological fitness (202). This drive underlies the remarkable curiosity we exhibit as a species, our hunger for knowledge, and our need for justification before adopting a belief. We are, in this light, cognitively predisposed to judge beliefs in terms of their truthfulness. Truth – or the concept of justified beliefs – however, remains vacuous without criteria in terms of which it can be realised or in terms of which these beliefs can be justified.

These criteria are epistemic values. They enable us to compare different representations and infer which one offers the best explanation - which representation, in other words, best approximates 'truth' and most justifies belief. According to Kuhn (1977:321-322), such values include: accuracy[30] (predicting all or most data and explaining away the rest), consistency (both internal and with other relevant and accepted theories), scope (the consequences of a theory should extend as much as possible beyond the data it is required to explain), simplicity (explaining the data as economically as possible) and fruitfulness (the degree to which a theory permits one to make new predictions).[31] More basically, they boil down to: predictive accuracy (a representation should be confirmed in its predictions by states of the world), coherence (the elements within a representation should not contradict each other, nor should the representation be in contradiction with other representations), scope (a representation should ideally explain all data) and, more controversially, a sense of aesthetics (between two theories explaining the same amount of data in a coherent way, the most elegant formulation – i.e. the most economical one – carries away our preference). Longino (1990:4) refers to these criteria as constitutive values. They are values 'by which to judge competing explanations' and are 'generated from an understanding of the goals of science'. They can be contrasted to contextual values, which are 'personal, social and cultural values'. The latter

[30] Kuhn uses the term 'accuracy' to point out the predictive accuracy of a representation. This is not to be confused with my use of the term, which – as pointed out earlier – refers to an appropriate relationship between representation and its object.

[31] Kuhn claims that these five criteria provide the shared basis for acceptance of one theory over another. However, he argues, this shared basis is not sufficient to determine scientific choice, i.e. to eradicate the incommensurability that governs over competing paradigms (Kuhn, 1977:331).

depend on the particular cultural context in which science is conducted, while the former are derived from the very enterprise of scientific enquiry itself.

Carruthers (2006:347) argues that those epistemic values – i.e. Longino's constitutive values – are most probably innate, for they seem universal to human cultures, from hunter-gatherer societies to western scientific communities. Those values underlie, in other words, not just modern scientific reasoning, but all of human kind's belief forming about the world. Furthermore, Carruthers points out, they are not – at least among hunter-gatherers – explicitly taught. Therefore, these epistemic values must be part of our cognitive endowment, much as our innate drive to search for truth.

I agree with this analysis on a more principled basis. How, indeed, would humans ever have transgressed their biological bias, if were they not already endowed with an epistemic value system inciting them to question their assumptions and look for alternative, *preferred* representations? Without these values or epistemic guidelines there could be no preference for one representation over another, and, without this preference, there could be no motivation for, or meaning in, producing alternative – i.e. transgressing – views. Therefore, just like our predisposition to look for truth, at least some of those values must be anchored in our innate cognitive make-up and cannot be purely cultural products, since culture itself depends on the transgression of our biologically constrained view of the world. Indeed, without the ability to transgress its innately predisposed ways of viewing the world, humanity would never have entered the cultural realm, in which the world comes to be viewed through a rich tapestry of diverse spatio-temporal perspectives instead of a singular species-specific view.

Furthermore, the claim that epistemic values have an innate basis is supported by empirical research on simplicity. Lombrozo's (2007:233-235) experiments point to a preference for simpler explanations (i.e. explanations invoking less causes) and the role of simplicity in probabilistic reasoning. Finally, it seems hard to conceive of the possibility that we are predisposed towards truth without possessing epistemic criteria. Indeed, how could we be endowed with an innate drive to represent the world truthfully, without the necessary tools to respond to this drive?

7. Conclusion

I set out to uncover the cognitive abilities that enable us to go beyond and against the view of the world we are predisposed to have in virtue of the particular architecture of our senses and cognitive wiring. I argued that this ability is not the result of either a lack or excess of innate predispositions, but of cognitive operations we perform on the species-specific representations yielded by our sensory and cognitive apparatuses. More precisely, there are three cognitive abilities that enable us to transgress our commonsense representations.

The first is our ability to metarepresent, representing the representations we hold and their parts. The second is our ability to produce alternative representations by recombining elements we gather from input and conceptual domains unlimitedly and transferring our modes of representing to other previously non-conceptualised domains. The third is our possession of an epistemic value system, enabling us to rank representations with regard to epistemic desirability.

These three abilities are the cornerstones of the human ability to transgress its biological bias, and are therefore the source of the unique cognitive achievements that characterise homo sapiens and distinguish it from all other creatures on this planet. They provide us with the cognitive flexibility and creativity that enables us to overcome the outputs of our 'hard-wired modules', parting with our nature into this awe-inspiring diversity of human culture.

This, of course, does not imply that all transgressing theories can be accounted for merely in terms of these three cognitive faculties. Indeed, quite obviously, someone living in 5000 B.C. could not have come up with the theory of relativity. To do this, he or she would have to have stumbled upon a number of crucial astronomic discoveries (not the least of which is that the earth is round and orbits the sun), developed mathematics to a breath-taking degree of complexity, and developed the proper technology for all these astronomical discoveries (at the very least, a powerful telescope). Einstein's ability to formulate his transgressing view of the universe, in this light, is not the just the product of his own cognitive abilities but – to a very important extent – of the impressive body of accumulated knowledge he was born into.

Indeed, theories do not originate from nothing – they are built upon previous theories, which again are founded on an older set of theories. However, the very possibility of embarking upon this chain of theories – transgressing its species-specific set of uncritical representations, as homo sapiens has done – is grounded in this three-fold ability to take representations as objects of cognitive attention, to produce alternative representations with the available resources, and to give these representations an epistemic orientation.

These operations enable us to produce an unlimited amount of representations, to divorce subject matters from the reasoning patterns we are predisposed to apply to them, and to think about subject matters for which we have no innate reasoning module at all. This opens up a radically different epistemic relation to the world other than the particular relation determined by our nature. It does, indeed, endow us with an 'open' epistemic relation, not yielding a fixed set of representations, but representations which are both unlimited in number and subject to change. In the following chapter, I will gauge the scope and limits of these transgressing theories, bringing to bear the question that fuelled this dissertation: what can we know about the world?

Chapter 4: Scope and limits of transgression

1. Introduction

In the previous chapter, I argued that we can transgress our biologically based views of the world. We can, in other words, represent the world in ways that go beyond and against the representations we are predisposed to have in virtue of our nature. In this chapter, I will gauge the scope and limits of these transgressing representations.

Arguing that we can both shape an infinite amount of possible representations and that we are not bound to represent any subject matter in a particular way, I characterise the cognitive relation we entertain with the world as an open cognitive relation. Indeed, while other organisms on this planet are forced to represent their environment in a particular set of ways, entertaining a closed cognitive relation with the world, we exhibit no such constraints.

This, however, does not imply that there are no limits to the representations the human mind can produce. From the fact that we entertain an open cognitive relation with the world, it does not follow that we can represent the world in any (theoretically) possible way. In this regard, we cannot represent any possible subject matter in any possible way, merely in an unlimited *amount* of ways within the limits imposed by our nature. Those limits, I argue, are set by the particular senses and cognitive reasoning patterns that evolution has provided us with.

Our perceptual and cognitive nature limits the ways we can represent the world either by biasing us against representing the world in a certain way, or by straightforwardly closing us off to certain possible representations. With regards to bias, I argue that due to the nature of our senses, some entities and properties of the world are easy to detect, while others are almost impossible. We are, in other words, biased against detecting some aspects of the world because of the particular nature of our senses. Furthermore, while we are able to transgress our commonsense representations, our intuitive grasp of the world still plays a role in the transgressing representations we produce. In this light, we are biased against representations which depart too radically from our intuitive understanding.

With regards to closure, I argue that when conceptualising about the physical world, we (necessarily) engage our perception based imagination. This entails that we are bound to model the world by means of the sensory resources at our disposal. Representations, therefore, which resist being grasped by those resources, remain forever beyond our ken. Furthermore, since we transgress our commonsense representations not by applying new reasoning patterns, but by recombining the perceptual and conceptual resources at our disposal (cf. Chapter 3 - 6.2 Variation through recombination), we are closed off to representations which cannot be deconstructed in the 'cognitive building blocks' that evolution has provided us with. We are, in other words, closed off to representations which cannot be reached through the particular cognitive reasoning patterns we possess.

2. Scope of transgression

2.1 Productivity of thought

One of the main features of natural languages is their combinatorial nature. The meaning of a sentence is derived from the meaning of its constituents – i.e. the words – and the grammatical rules underlying the ways in which these constituents are combined. This enables us to understand sentences we've never encountered before, by grasping the meaning of the words and the combinatorial rules. Consider, for example, the sentence: 'all young goldfish like the sound of the wind blowing when staring at the moon'. This sentence has most probably never been encountered before by the reader, but this doesn't prevent him or her from understanding it. Indeed, the only condition for something to be intelligible is the ability to understand the parts and the grammatical construction. This allows us to conceive of and understand an endless series of new meaningful expressions – i.e. new thoughts – with finite means – i.e. the memorisation of a finite set of words and rules.

As Pinker (2002:37) points out, even with a few thousand nouns and verbs, there are already several million correct grammatical ways to open up a sentence. Add prepositions, adverbs and adjectives to that, and the combinations quickly multiply to an astronomical number. Moreover, because of the possibility of recursion – phrases containing other phrases – the number of possible, grammatically correct sentences (and therefore meaningful sentences) is in principle infinite. Consider the phrase: 'I heard that you think that she fears that he knows that …'.

From this productivity of language, yielding an endless possibility of different sentences, we can derive the productivity of thought. Since every meaningful sentence expresses a thought, there must be an infinite amount of different possible thoughts the human mind can entertain. We can, in this regard, produce an infinite amount of distinct and meaningful representations of the world.

2.2 Representational flexibility

Another characteristic feature of human cognition is – as discussed at length in the previous chapter – our ability to represent subject matters differently from the way in which we are predisposed to represent them. Indeed, while any other organism on this planet will represent its surroundings in a particular species-specific way, showing no, or very limited, flexibility, human beings come up with a rich and inexhaustible variety of conjectures in explaining subject matters.

Greek astronomers overthrew our spontaneous assumption that the earth is flat; Copernicus's heliocentric model placed the sun in the centre of the universe, affirming that the earth is in constant motion; and Einstein's theory of relativity contradicts our deeply rooted intuition that time and space are two independent realms, predicting that time is dependent on velocity. Moreover, atomism postulates that matter is constituted out of tiny particles floating around in ether, and string theory conjectures that the electrons and quarks making up these atoms are one-dimensional oscillating lines.

This intellectual creativity that humankind displays in conceptualising the world is not, however, restricted to scientific theories alone. Since the dawn of humanity, societies have produced an astonishing variety of mythological stories accounting for the origin of the universe. In ancient Greece, Gaia (the earth) was said to have emerged from chaos, giving birth to the sky. Cherokee Native Americans believe that the earth was created when a water beetle coming from the sky dived to the bottom of the endless pool of water bringing up some soft mud to the surface. Taoism, on the other hand, conceptualises the origin of the universe as the breaking of a cosmic egg, releasing sky and earth, which both expanded, growing further apart, and the African Bantu tribes invoke a vomiting God (Leeming, D & Leeming, M, 2009). Moreover, different cultures represent humanity and its environment in different ways. Hinduism tells us that all living creatures are temporary vehicles for souls caught in a cycle of reincarnations. Animistic religions, on the other hand, claim that the whole world is imbued with soul, while the Judeo-Christian tradition reserves the possession of a soul to humankind alone. Furthermore, genesis tells us that humans are created in the image of God, while Darwinism postulates that we are but a particular species in the primate family. Finally, some philosophers have claimed that nothing exists but in the eye of the beholder and that we might well be stimulated brains in vats. In short, there seem to be no limits as to how we come to conceptualise the world, ourselves and everything else in it.

This inexhaustible creativity is the true hallmark of human cognition. Our ability to transgress our predisposed views of the world opened up an endless variety of possible representations. Cutting loose from biological determinism, the human mind entered the kaleidoscopic realm of culture, representing the world in a myriad of different (and contingent) fashions. Indeed, while we share species-specific knowledge systems that we are inclined to apply to our environment, we obviously go beyond (i.e. we think about different subject matters than the ones we are predisposed to think about) and even against these predispositions (i.e. we contradict our universal commonsense representations).

2.3 Open cognitive relation

The fact that we can produce an infinite amount of meaningful representations (cf. productivity of thought) and are not bound to represent any given subject matter in a particular way (cf. flexibility of representation), endows us with a different cognitive relation to the world from that of any other organism on this planet. Indeed, since all our intuitive representations can in principle be rejected and substituted by an infinite amount of possible representations, the human mind is not forced to represent the world in a particular way, or even a particular set of ways. It is, in other words, not bound to view the world according to any of its intuitive core beliefs, and is able to generate an infinite amount of theories about the world. The cognitive relation humankind entertains with the world, therefore, can be characterised as 'an open cognitive relation', as opposed to the closed cognitive relation other species on this planet entertain with the world.

An open cognitive relation entails that every belief about the world can in principle be rejected and replaced by another representation. There is, in this regard, a contingent relation between representing subject and represented object. The objects of representations are, in other words, not necessarily represented in a particular way, nor are the ways in which those objects can be represented (i.e. the amount of possible representations that can be generated) exhaustible. This makes human representations of the world radically unpredictable. There is no way of deducing the representation that will be shaped in human minds, based upon the object it represents and its cognitive apparatus.

Indeed, while a gosling will invariably represent the first moving object it perceives after hatching the egg as its mother (Lorenz, 1935), or a tick an object of 37° C which exudes buteric acid as a food source (Uexküll, 1909), human beings can represent their environment in an unlimited number of contingent ways. Humans both represent their environment as the expression of metaphysical forces, or as a closed causal chain of events governed by universal laws. They attribute supernatural essences to the natural world – as is the case in animism – or reduce all life to lifeless atomic interaction. They can even view themselves as either radically different creatures from all other organisms on this planet – rational beings, as Aristotle has it – or as merely another species within the primate genus. Finally, humans both explain the origin of the world as the intentional creation of a divine being or as the result of the sudden expansion of an extremely dense 'singularity'.

This, however, does not mean that there is no common ground to be found when comparing the ways in which different cultures represent the world. As pointed out in chapter 1, comparative anthropology and developmental psychology reveal universal cognitive predispositions to represent aspects of the world in particular ways (cf. folk physics and folk biology). It only means that, while these cognitive universals underlie all human thought, they do not produce an invariable set of representations shared by all human beings, and this is precisely because the human mind is able to transgress the commonsense representations it holds in virtue of these universals.

In this light, human culture cannot be reduced to human nature, nor can it be completely dissociated from it. This is at the basis of Wilson and Lumsden's (1981) 'leash principle'. Human culture, in all its mind-blowing diversity, ultimately remains connected to human nature. Similarly, Ruse (1995:158) claims that culture 'sits on top of a bed of biological constraints and dispositions' - it is 'the flesh which adheres to the skeleton of biology'. This, however, as Ruse points out, does not imply that culture is completely controlled by biological forces (158). Indeed, while our biology undoubtedly underlies the ways in which we come to view the world, it does not straightforwardly determine it. Culture, in other words, wanders 'freely' within the boundaries set by nature.

2.4 Conceptual spaces: reconciling openness and closure

The fact that we entertain an open cognitive relation with the world does not entail that there are no limits to the ways we can represent the world. As pointed out, nature still imposes boundaries on our cultural expressions. Human thought, therefore, while able to produce an infinite amount of representations of the world, is still constrained by both the particular perceptual and cognitive abilities producing these representations. We can, in other words, not represent any possible subject matter in any possible way, merely in an infinite *amount* of ways within the limits imposed by our nature.

Boden's (1990:89-91) concept of a 'conceptual space' is enlightening in this context. Drawing from research in Artificial Intelligence, Boden defines a conceptual space as the space of computational possibilities a system can generate. Such a 'generative system', as Boden puts it, is composed out of data and action rules, enabling the system to generate a number of locations. Chess, for example, allows for a number of possible board-positions based on the particular data – i.e. the different pieces and the structure of the board – and action rules – the rules by which those pawns can be moved across the board. Based on these data and rules, all possible board positions can be generated.

The number of locations within a conceptual space is dependent on the data and rules of the generative system. The possible locations that can be yielded by a game of noughts and crosses are obviously fewer than those of a game of chess, where the total amount of locations reaches astronomical proportions. Other generative systems even yield an infinite amount of possible locations within their particular conceptual space. As pointed out, natural languages can express an infinite amount of intelligible sentences based on a finite set of words (i.e. data) and grammatical rules (i.e. action rules).

Conceptual spaces, therefore, can both contain an infinite amount of locations, and at the same time be contained within particular limits. In much the same way as there are an infinite number of cardinal values between the number two and three, there can be an infinite number of locations in a particular conceptual space. This is the case for human thought. Endowed with an open cognitive relation to the world, we can produce (in principle) an infinite amount of representations of the world, within the limits set by our perceptual and cognitive apparatus.

When exploring its conceptual space, the human mind does not visit locations in a random fashion. It is predisposed to apply certain reasoning patterns (action rules) to certain subject matters (data) and, therefore, is drawn towards certain locations within the conceptual space, while ignoring other locations. Just as a chess player will use certain guidelines – such as, for instance, 'protect the queen' – and will not try any permissible move at random, our mind does not wander blindly in its conceptual space, but selectively, yielding particular locations.

Boden (1990:89) refers to these guidelines as heuristics. They enable any cognitive system to move 'insightfully' through the conceptual space generated by the system. Heuristics, in Boden's words, 'prune the search tree' (91). They save the problem solver from having to visit every location in the conceptual space, by ignoring parts of it. These heuristics are necessary when you are dealing with an extensive and, a fortiori, an infinite amount of locations in a conceptual space, as is the case for the space containing human thought processes. While a successful game of noughts and crosses might still be played by selecting the proper move through testing every possible move by trial and error, this becomes impossible when the set of locations within a conceptual spaces is increased exponentially. Indeed, while a computer with enormous computational power might still use trial and error in chess, processing astronomical numbers of potential moves per second, the human mind is forced to 'prune the search tree' and to use pathways in exploring the conceptual space.

In this regard, when it comes to representing our environment, natural selection has provided us with a number of heuristics. These heuristics yield a particular set of locations – referred to by Carey and Spelke (1994) as innate knowledge systems, by Mithen (1996) as domain-specific knowledge systems, and by Carruthers (2006) as modules – which determine the way we are predisposed to conceptualise about particular aspects of the world (cf. chapter 3 – 4. Mapping across domains). Without these heuristics, it would be virtually impossible to stumble upon useful representations – i.e. representations which boost an organism's chances of survival and reproduction – leaving the organism at the mercy of blind trial and error. In other words, natural selection has provided us with pathways within our conceptual space, zoning in on locations that boost our fitness, while ignoring other parts.

However, as pointed out, we can transgress those innately based representations – i.e. we can explore different locations in our conceptual space other than those which natural selection has led us to. The human mind has, indeed, the distinctive ability to forge new pathways in its conceptual space, exploring new locations and therefore representing the world in ways that go beyond and against its predisposed views. Transgressing representations, in this regard, are the result of changing heuristics. They can't, however, take us beyond the limits of our conceptual space. A conceptual space, as pointed out, is defined by a generative system that comprises a particular set of data and a particular set of action-rules. Together they yield all the possible locations a cognitive system can generate. In this light, we are endowed with a particular set of data – i.e. the input we gather from the world – and a particular set of action-rules – i.e. the cognitive operations we can perform on these data. Based on this, we can generate a particular set of representations. The fact, however, that these representations are not limited in number, does not – as pointed out above – entail that our conceptual space contains all possible locations.

In the next section, I will investigate the ways in which our sensory and cognitive apparatuses limit the transgressing representations that the human mind can yield. Drawing from this view of human cognition as a generative system, defining a conceptual space and endowed with heuristics to explore that space, I will argue that we are both biased against producing some representations, and straightforwardly closed off to the possibility of shaping some representations. A bias exists when, due to the nature of our perceptual and cognitive abilities, we are unlikely to reach certain locations within our conceptual space, while closure refers to the locations outside our conceptual space.

3. Limits of transgression

While we can transgress our commonsense representations, endowing us with an open cognitive relation to the world, our sensory and cognitive faculties do, nevertheless, restrict the possible representations we can produce. Indeed, as pointed out above, we can both be biased against representing the world in a certain way due to the particular nature of our senses and cognitive apparatus, and closed off to certain representations. In terms of Boden's (1990) conceptual space, the representations we are biased against shaping remain within our conceptual space, but require pathways we are unlikely to employ, while the representations that we are closed off from shaping are located outside the limits of our conceptual space.

3.1 Bias from representing the world in a particular way

3.1.1 Sensory input

We represent the world based on the input we gather from it. Through the senses that we have evolved, we receive a particular set of input from the world. These sensory data are the objects of our representations. When we see a stone falling to the ground, we explain this by referring to its innate tendency to be on the ground, or by referring to the law of gravity, according to which a greater mass attracts a smaller one. Similarly, we explain thunder and lightning in terms of the intentional action of an angry god, or as the electrical discharge from colliding clouds. Every theory about the world, in this light, is constructed from the data we gather from the world.

However, as pointed out in chapter 1 (cf. 2.4.1 Perceptual closure), we only perceive particular aspects of the world. Only a small part of the available stimuli trigger our senses. Indeed, we perceive stimuli only within particular ranges, and endowed only with a certain level of resolution; and, most importantly, there are vast realms of potential data in the world for which we simply have not evolved the appropriate sensory receptors. This leaves us with a narrow scope on the world.

Nevertheless, we extend the reach of our senses through 'artificial detecting devices', allowing us to gather input from the world that falls beyond the scope of our naked senses. We create telescopes, microscopes, antennas and stethoscopes, enabling us to detect otherwise unperceivable entities both in the macro- and microscopic realm. Furthermore, rather than just increasing the resolution of our senses, we possess devices that detect ranges of phenomena we cannot perceive, such as, for instance, the invisible part of the electromagnetic spectrum, and the inaudible frequencies of sound. Finally, we even detect phenomena for which we have no sensory receptors at all, such as air pressure, through barometers and magnetic fields.

In this regard, assuming that the physical world is causally closed – i.e. that all entities exert a causal effect on other entities – we could, in principle, detect every entity and property in the world. Indeed, while we might not be able to observe some aspects of reality, if these aspects influence other aspects, which again change others ad infinitum, all elements in the world must eventually leave a trace that we can gather through observation and its mechanical extensions. Physical entities do not, indeed, have to be detected by a particular sense 'designed' to detect it. Light, for instance, can yield auditory stimuli, just as sound can be translated into graphs by a computer. In this context, interestingly enough, the theory of the big bang was first confirmed by the accidental recording of a persistent low noise by a six meter horn antenna originally built to detect radio waves bounced off echo balloon satellites. This mysterious sound turned out to be caused by radiation released by the big bang (Dycke, et al, 1965). In other words, light wavelengths originating from the blast at the birth of the universe interfered with this hyper sensitive antenna, causing it to produce a faint, constant noise. In this regard, all physical phenomena could, in principle, be detected, as long as they exert even the tiniest influence in the causal chain of physical events.

Nevertheless, some entities or properties of the world are easier to detect than others. The existence of the sun is much easier to detect than the existence of Jupiter's moons, since we are able to detect the former with the naked eye and the latter only with telescopes. Similarly, the existence of Jupiter's moons is easier to detect than the existence of the above mentioned radiation caused by the big bang, since the latter cannot be detected by a 'simple' telescope, but requires high-tech radar equipment and the proper scientific hypothesis – i.e. the birth of the universe in a big bang – to explain the particular recording. Therefore, while it might be the case that no physical entities are impossible to detect in principle, the detection of some entities is rendered much easier than the detection of others by the nature of our particular senses.

In this context, imagine that human beings did not possess the sense of vision. All other things being equal, they could feel, smell, taste and hear, as well as apply all our cognitive abilities to this input. Their blind scientists, it seems, would have to overcome enormous odds to detect some of the physical entities and properties we observe without much trouble through our sense of vision and its extensions. Indeed, being aware of shapes by touch, they could come up with a similar spatial representation (objects dispersed in space). Assuming that they invented some orientation device based on magnetism, they could, by crossing the earth in straight lines, even come to represent the planet earth as round. This discovery would lead them to conceptualise their planet as a round sphere surrounded by ether. Furthermore, they would be aware of a regular alternation of day and night by the obvious difference in temperature, and could conceive of a heat source that crosses the surface of the earth during a certain period of time from one side to the other.

However, the conceptualising of this heat source as a celestial body at some distance from the earth, and the movement of this heat source as the result of the spinning motion described by the earth, would be anything but evident. Furthermore, knowledge of the existence of the other planets in our galaxy, let alone the existence of other galaxies and stars, would have been almost impossible to achieve without vision (and its extensions), which offers all these entities in plain sight. The discovery of planets (literally: wandering stars), and later, the heliocentric model of the universe, are, in this context, the product of careful stargazing. Without vision, it bears no doubt that scientists would have to overcome enormous odds to zone in on these basic insights into our universe. They would need to capture all this input from the universe, which is currently simply given to us at the end of a telescope or even by plain sight, by other means. While light radiation can, as pointed out above, yield an auditory signal, it seems, however, a Herculean task to unravel the layout of our galaxy and other nearby galaxies from mere auditory signals.

Furthermore, awareness of the existence of microscopic entities seems equally problematic. Undoubtedly, our blind scientists would have enormous difficulties in detecting the existence of microscopic organisms or the cells that make up all organic bodies on this planet. While they could have basic ecological and anatomical knowledge, representing the interaction of organisms on this planet and mapping organs and tracing their function, it would be extremely difficult for those blind biologists and anatomists to form an accurate and precise representation of either microscopic organisms (as microbes) or the microscopic structure of organic tissue – i.e. cells. They might make conjectures of such entities fulfilling the function that they do, much as we conjecture about atoms, electrons and quarks without observing them, but without receiving any sensory input from these entities, it is very unlikely that these hypotheses would represent these microscopic entities as precisely.

Finally, distinctions we draw in our environment based on visual data alone would obviously be problematic to make. This could lead their zoologists to mistakenly group different kinds of butterflies (which only differ in the patterns on their wings) under the same species, their botanists to fail to distinguish red from white roses (granted they only differ in colour), and their evolutionists to ignore the adaptive feature of camouflage in animals. Similarly, their geneticists would be much more likely to overlook the heredity of hair and eye colour, and their ethologists the cause of some animal behaviour, granted that it is triggered by vision alone.

In this regard, due to the nature of their sensory input, our blind scientists would be biased against discovering the existence of some physical entities and properties. They would, in other words, have enormous difficulties in discovering elements we gather by merely opening our eyes. Therefore, while we are not restricted to the data we access through our naked senses and can – in principle – extend our perceptual reach to possibly all physical entities and properties, the detection of some becomes an extremely difficult task because of the particular nature of our senses. Indeed, we could very well be those blind scientists for some other species which has more senses and can therefore access entities without any problem for which we would need a whole arsenal of complex technological support.

Detecting mechanisms, in this sense, are better described as extension devices of existing senses than radically new and different artificial senses. In discovering data from the physical world, we start from what we gather through our naked senses, and extend our reach from there. We do not delve into completely different realms of data with these artificial mechanisms at once, but gradually work our way up to more and more inaccessible data. Indeed, how could we construct a mechanism that measures something of which we are totally ignorant? We can only direct our gaze to entities we know or at least have reason to suppose exist. This ‘sharpened gaze’ through mechanical detection devices then opens up the existence of new entities, which again enables us to focus our gaze, providing us with a new set of bearings to further direct our detection of aspects of the world.

Therefore, while we might not be irremediably closed off to some physical entities and properties, we are still biased against detecting some aspects of the world because of the particular nature of our senses. This entails that some representations, requiring the detection of such aspects, could very well be nearly impossible to reach, like macro- and microscopic structures to our blind scientists. Indeed, we extend our reach to unobservable entities, transgressing the subset of input data our senses provide us with, but can only do so by gradually departing from the data that are already given to us. This provides us with a perceptual bias: some elements will be (relatively) easy to discover, because of their proximity to elements we already observe, while others might be extremely hard for us to gather. More or different senses, in this regard, would provide us with a different set of input that we observe through our senses and, therefore, a different perceptual bias, making different elements of the world more accessible.

3.1.2 Intuitive grasp

While we can transgress our intuitive or commonsense representations of the world (cf. chapter 3), these intuitive grasps of the world still underlie to an important extent the way we come to represent the world in our scientific – i.e. transgressing – attempts to explain the world. As De Cruz and De Smedt (2007) point out, although science has parted ways with intuitively based folk theories, we can nevertheless expect that they will continue to play a role in scientific enquiry, since the human mind is evolved to understand objects in the world according to these intuitive categories.

Considering the case of scientific theories about human evolution, De Cruz and De Smedt (2007) argue that not only do intuitive ontologies[32] (i.e. the intuitive categories and their respective modes of inferences underlying our folk theories) shape intuitions about human evolution, but they also guide the direction and topics of interest in the research programmes. In exploring the relationship between intuitive ontologies and the scientific discourse surrounding human evolution, De Cruz and De Smedt point to two distinct intuitive modes of understanding that are relevant to the issue of human evolution; namely, the human-nonhuman distinction and psychological essentialism. The first is a psychological mechanism enabling us to distinguish conspecifics from nonconspecifics, an adaptation we share with most species, helping us to, among other things, recognise potential mates. The second, on the other hand, makes it possible to override perceptual differences, enabling us to make inductive inferences about food, predators and other ecological features relevant for survival and reproduction. While the first intuitive ontology – i.e. the human-nonhuman distinction – leads us to consider human evolution as exceptional, the second – i.e. psychological essentialism – leads us to derive from the extensive similarities between humans and apes that

[32] Boyer (2000:277), from whom De Cruz and De Smedt borrowed the term 'intuitive ontologies', describes it as 'a series of category-specific intuitive principles that constitute an evolved 'natural metaphysics'.'

they share the same essence, minimising the difference in psychological abilities that both species are endowed with (358).

The tacit assumption that the human species is unique has led palaeoanthropologists to a unilineal view of human evolution. Mayr (1950:115-116), for instance, argued that all hominids can be grouped in a single lineage from australopithecines through homo habilis to homo sapiens. This, however, is in sharp contrast with the usual branching pattern of evolution, and was explained away by Mayr by invoking the fact that hominids did not speciate because, possessing culture, they occupied more ecological niches than any other species (De Cruz and De Smedt, 2007:359). In other words, because they entered the cultural realm, human beings are part of a unique evolutionary process. Recently, however, the finding of homo floresiensis, a small hominid with the brain size of an australopithecine, dated at 18 000 years B.C., completely overthrew this unilineal view on human evolution. Nevertheless, despite the discovery of many more genera within the hominid linea in the last decades, palaeoanthropologists still attempt to prune the tree of human evolution (De Cruz and De Smedt, 2007:361).

Essentialism, on the other hand, leads us to minimise the difference between humankind and its evolutionary cousins. As De Cruz and De Smedt (2007) point out, from a radical separation between human and nonhuman, we often succumb to the opposite temptation of eradicating the difference altogether. In this context, Diamond (1992) argues that an extraterrestrial observer would objectively classify humankind as the 'third chimpanzee', thereby overlooking the blatant fact that this third chimpanzee writes about the two others in fluent grammatical language (De Cruz and De Smedt, 2007:363)! Moreover, much of current research in comparative psychology is directed at finding similarities between humans and nonhuman primates. Not surprisingly, the interpretation of experiments aimed at gauging primate possession of a theory of mind, for instance, often betrays a bias towards anthropomorphising the apes (364).

Furthermore, more than just directing and biasing scientific research, intuitive ontologies also exert a limiting function upon the extent to which they can be transgressed. We are, indeed, biased against representations which depart (too) radically from our predisposed views of the world. Such counter-intuitive hypotheses about the world often fail to convince us, and incite us to produce representations which are more in line with our intuitive understanding of the world. As De Cruz and De Smedt (2007) point out, without intuitive notions to guide them, scientists often fail to agree on even the most basic foundations of their field. In this regard, evolutionary biologists still disagree over what constitutes the basic unit of selection, what constitutes a species and how evolution takes place (i.e. as a gradual process or one of punctuated equilibriums) (357). Even more so, in physics, research fields such as quantum theory and M-theory, better known as string theory, remain highly speculative, and their distinctively counter-intuitive implications cause wide-spread criticism and disagreement. In these cases, without a conceptual grip provided by our intuitive notions, we are doomed to remain cognitively unsatisfied. Indeed, when these speculative conjectures violate our sense of causality or our basic intuitions about space and time (stating, for instance, that a particle can both be in space A and B at the same time T) – as is the case in quantum theory – we cannot easily accept them. This leads to a situation where science is either in line with intuitive ontologies, and can therefore be suspected to be biased, or departs radically from these ontologies, but strikes us as distinctly unrealistic.

In this regard, while we are able to transgress our commonsense representations, the intuitive ontologies at the core of these representations still play a role in the transgressing representations we produce. As shown by De Cruz and De Smedt (2007), intuitive ontologies are pervasive in scientific discourse, directing research and biasing the interpretation of experimental results. When science, on the other hands, departs radically from our intuitive grasp of the world, violating some core intuitions, we cannot invest the same amount of belief in these representations and find our minds turning away from these conjectures, looking for alternatives which are more in agreement with our guiding intuitions.

Therefore, much as our senses bias us against detecting certain elements of the world, our cognitive nature biases us against interpreting the world in certain ways. Transgressing representations, in this context, start from intuition and extend from there. The further away we move from our intuitive grasp of the world, the fewer guiding notions we have, and the more difficulty we have in accepting those representations. In this context, just as we can transgress the particular data yielded by our sensory apparatuses, but are still biased towards discovering data in the proximity of these data, we can transgress our commonsense representations of the world, but are biased towards producing theories which are still centred around these particular intuitive notions. Some representations, therefore, either because they require input data we are not likely to detect, or because they require us to overthrow our intuitive grasp of the world too radically, are situated at the outer limits of our conceptual space. While we are not closed off to those representations in principle, the pathways they require through our conceptual space in order to be reached are unlikely to be employed because of the working of our senses and our cognitive predispositions. There are, however, representations which are not merely unlikely to be shaped, but which are straightforwardly impossible to be conceived by the human mind. These are the locations situated outside our conceptual space and the subject of our next section.

3.2 Closure from representing the world in a particular way

3.2.1 Sensory resources

When conceptualising about physical entities, we engage our sensory based imagination. We reduce the physical world to basic elements such as water, air, earth and fire, we conceptualise it as atoms hovering in ether, and we even conjecture that those basic elements are waves or strings rather than particles. However far those models take us away from the way the world appears to us in perception, they nevertheless still yield to a sensory grasp. In this sense, when conceptualising sound as propagated waves and light as electromagnetic wavelengths – thereby transgressing the commonsensical confusion of the physical cause of these perceptions with their phenomenal qualities – we do so by reconceptualising them by means of the sensory resources at our disposal.

McGinn (1989:358) argues that the theoretical concepts we use to describe the physical world are formed by an analogical extension of what we observe. He refers to this as a 'principle of homogeneity' operating in our introduction of theoretical concepts on the basis of observation. This, according to McGinn, is at the basis of the insolubility of the mind-body problem, since inference to the best explanation of physical data (i.e. physiological data of the brain), will never take us outside the physical realm (i.e. the realm that can be conceptualised in sensory terms), leaving consciousness out of the equation. The core of the mind-body problem, in other words, according to McGinn, is that theoretical concepts about brain states will always be perception based or 'imagible' as Krellenstein (1995:235) puts it, and therefore unable to connect with consciousness, i.e. unobservable subjective experiences.

In this sense, McGinn (1989:358) argues, we arrive at the concept of a molecule based on our perceptual representations of macroscopic objects, conceiving of smaller scale objects of the same kind. Unobservable material objects, in other words, are conceptualised based upon observation based imagination. Such a theory of concept formation, as Krellenstein (1995:242) points out, does not pertain to all abstract concepts, but only to concepts providing causal explanations of physical entities and properties. Numbers or numerical relations, for instance, are not grasped by means of perception based models.

McGinn's case for this homogeneity constraint, however, is subject to critique. Kukla (2005:69) argues that this principle of homogeneity can either be considered normative – i.e. theories should conform to this principle – or descriptive – i.e. our mental apparatus *cannot* explain the world without this analogical extension of the senses. The interesting case here, and the one McGinn needs to prove, is the latter. This descriptive principle, according to Kukla, can 'easily' be shown to be false. It is, he continues, disproved by the intentional concepts of folk psychology, on the one hand, and by the concepts of quantum mechanics, on the other hand. In folk psychology, as he points out, we introduce intentions, desires and so on, based on our observation of human behaviour. These intentions cannot, according to Kukla, be seen as analogical extensions of perceptual concepts (69-70). This counter argument, however, does not concern us, since I only made the case for a perception based introduction of theoretical concepts pertaining to physical entities and properties.

The second counter argument, however, needs answering. The quantum mechanical concept of superposition of states can, according to Kukla (2005), hardly be regarded as an analogical extension of observational concepts. While this can still seen to be the case for classical conceptions of atoms and electrons, Kukla argues, it no longer holds for the quantum-mechanical postulation of the ubiquitus electron, as it is not at one location at a specific time. This according to Kukla, overthrows the descriptive principle of homogeneity, since we can not imagine an electron being at more than one location at a time, yet we are still able to postulate it (70).

With the replacement of Bohr's atom as a miniature solar system by Heisenberg's quantum mechanics, in which particles could not be visualised because they have purely mathematical meaning, the observational basis (or imaginability) of the physical theory describing sub-atomic interactions was indeed lost. This, however, as Miller (2000) points out, does not mean that the role of visual imagination in modelling the physical world came to an end. As he points out, Heisenberg, and later Feynman, strove towards a new form of visual imagination, resulting in the Feynman diagrams. These diagrams, rather than reflecting the actual physical processes, enable scientists to think imaginatively about the underlying mathematics. In other words, rather than reflecting the structure of the physical world – which as Kukla (2005) rightly points out, can no longer be visually reflected – they reflect the structure of the theory. Miller (2000), in this regard, distinguishes between visualisation and visualisability. While visualisation refers to visual imagery which is abstracted from phenomena given in perception (e.g. Bohr's atom as a miniature solar system), enabling scientists to imagine outcomes in their 'mind's eye', visualisability refers to this new form of visual imagination, where the scientist no longer attempts to visualise sub-atomic entities but nevertheless uses visual imagery to represent the mathematical structure of the theory.

Furthermore, quantum-mechanics still conceptualises the interaction of matter and energy as both particle and wave-like behaviour, thereby using visually graspable concepts. In this light, the visual factor remains present in this 'extremely transgressing' theory, both as a support for scientists to grasp the mathematical structure of the theory (cf. Miller's visualisability), and, more importantly, in order to conceptualise the physical entities themselves – i.e. the sub-atomic entities viewed as both particle and wave-like.

In this sense, there are different levels of perception based representations. On a commonsense level, we equate what we perceive with what is. The second level transgresses these representations by reconceptualising physical entities and properties with their perception based imagination. The third and final level throws this imaginability overboard, theorising about physical entities in such a way that it can no longer be visualised. This level, nonetheless, does not do away with the perceptual base of the model as such, but merely with the possibility of playing it out in our 'mind's eye'. As Krellenstein (1995) points out, the more perceptually remote and – he argues – therefore less satisfying concepts, must still be 'imagable', i.e. they must still be perception based (albeit not yielding to visualisation) (247). There is, in other words, a bias – a decreasing cognitive satisfaction the more perceptually remote the model becomes – ending in closure, that all models must still be 'imagable' – they must still draw on perceptual resources.

Indeed, we cannot reduce our models of the physical world to mathematical equations or other non perceptual models alone. Barbour (1974:29-30), in this regard, distinguishes between different kinds of scientific models. The first are experimental models, designed to solve practical problems such as, for instance, wind-tunnels that are used to gauge the lifting force of a particular wing structure of an airplane. The second, at the opposite extreme, are logical models – formal deductive systems based on axioms and theorems. Those models deal entirely in the realm of ideas. The third kind of model is what Barbour calls mathematical models. They are symbolic representations of quantitative variables in physical and social systems, such as, for instance, equations expressing the relation between supply and demand in economics. Those models mirror their object in formal structure. The fourth kind, finally, are theoretical models. Rather than just enabling us to make predictions by representing quantitative variables, those modules are aimed at understanding their object. Their intent is to represent the underlying structure of the world. In order to do so, they postulate imaginative mental constructs accounting for the observed phenomena.

According to Barbour (1974), these mental constructs are shaped by analogy with familiar mechanisms and processes. Theoretical models, such as, for instance the billiard-ball model of gas, which postulates that gas is composed of tiny spheres bouncing around like colliding billiard-balls, then enable the development of theories, involving equations interrelating (in the given example) mass, velocity, energy and momentum of these hypothetical spheres (30-31). These models, Barbour argues, need not be picturable as such. We can, indeed, selectively suppress visual features, such as when imagining colourless elastic spheres. Nevertheless, they must be conceivable. They must, according to Barbour, be intelligible as units, providing us with 'a mental picture whose unity can be more readily understood than that of a set of abstract equations'. They do, therefore, make use of visual imagery (33).

Leaving this perceptual basis completely behind and reducing the physical world to mathematical equations, in this light, would both rob us of the framework necessary to develop theory, and – while still enabling us to make predictions (cf. mathematical models) – render us mute as to the nature of the objects of our study. Indeed, as Barbour (1974) argues, theoretical models in science account for observed phenomena, not just by predicting them, but also by explaining them: 'representing the underlying structures of the world' (30). Their subject is the (physical) object or phenomenon itself, not just the quantitative variables those phenomena are governed by. Theoretical models, in this sense, cannot be reduced to purely formal or mathematical models. According to Barbour, in opposition to an instrumentalist view of science, it is now increasingly accepted that explanation cannot be equated with prediction (41). Toulmin (1961), in this regard, cites the fact that Babylonians could predict eclipses from time tables, but could offer no reasons for this occurrence. Not possessing a heliocentric view of the universe, one can hardly say they understood the phenomenon.

Moreover, as Barbour (1974) points out, scientists report that visual imagery often predominates over verbal and mathematical thinking in scientific discovery (33-34). Visual models, in other words, drive the development of new theories. Furthermore, different (perception based) models yield different theories, defining different rules of correspondence between a set of postulated terms and observable variables (31). Our grasp of the world, therefore, is dependent on the perception based models we have at our disposal. This entails that we would be closed to representations of the world which resist being grasped through the perception based models we can shape.

Those models are the product of our sensory apparatus. The fact that mental imagery is connected with our sense of vision is – next to being supported by our commonsense observation – confirmed by extensive empirical evidence (cf. Kosslyn, 1980). Therefore, the particular senses we have evolved provide grounding to the perception based models we produce in theorising about the world. This entails that different senses would have provided us with a different substrate to model the world, and would therefore have yielded different representations of the world, which are now irremediably beyond our ken.

Moreover, rather than being restricted to our five senses in modelling physical entities and properties, we are restricted to vision alone. Evolution did, indeed, provide us with a dominant sense. As pointed out in chapter 1 (cf. 4.1.3 Human perception), as primates we rely primarily on our sense of vision. This entails that our representation of the world is mainly a visual one, as opposed to many other mammal species who rely more on their olfactory and/or auditory sense. This dominant sense, I argue, rather than just providing us with the set of data upon which we rely the most in our interaction with the environment, also underlies the way in which we – as cognitively highly developed primates – come to conceptualise the physical world.

Indeed, as shown above, all physical entities are conceptualised in a visual way by human beings. The physical cause of sound is viewed as waves, and that of taste and smell is viewed as particular molecules, which – like all molecules and sub-molecular parts, for that matter – are conceptualised visually. Similarly, heat is explained in terms of the movement of atoms, as are the three states of solid, liquid and gas. Finally, matter itself is conceptualised in terms of molecules consisting of atoms which consist of electrons circling neutrons and protons, which again are conceptualised as quarks, etc. Whether matter is ultimately reduced to particles or waves or even strings, we always model it in a visual way.

Evidence for this role of vision in cognition can be found in language. As pointed out in the previous chapter, language can be seen as a mirror of thought (cf. chapter 3 – 4.2 Analogy and metaphor), externalising our thought processes. In this context, interestingly enough, natural languages boost an impressive set of visual metaphors for cognition. We exclaim ‘I see’ to indicate that we understand something, ask someone to ‘clarify’ a concept, or point out that something is ‘clear’. Moreover, when having a new understanding of something, we refer to this as an ‘insight’, and when we can’t remember something, we have a blank. Finally, somebody’s opinion is often phrased as a ‘view’, and a great idea as a ‘vision’. Seeing and thinking seem to be inextricably linked for homo sapiens, and when conceptualising about the (physical) world, the limits of our models are set by the limits of our vision based imagination. In this light, it is no coincidence that even the very concept of 'imagination' is a visual metaphor.

Therefore, if evolution provided us with different senses or even just a different dominant sense, we would have been endowed with a radically different substrate for these conceptual activities. We would, in other words, represent the world with a radically different set of conceptual tools, not confining our models to visual imagery (as particles, wavelengths or strings) but to a whole different set of modelling tools. The fact that it is impossible for us to think about physical entities and properties outside this visual framework shows that we are irremediably confined in our conceptual imagination to use visually graspable data. How, indeed, could we make sense of atoms, electrons or quarks with olfactory or auditory models? Considering that our grasp of the world depends on the modelling tools at our disposal (cf. Barbour, 1974), there is no escaping the conclusion that we are irremediably closed to representations of the world which resist being grasped through the visually based models we must apply to physical entities and properties.

3.2.2 Cognitive resources

As pointed out in the previous chapter (cf. 6.2 Variation through recombination), we can only transgress our predisposed views of the world by recombining the elements we gather from our perceptual input and the cognitive tools at our disposal. Creativity, as Boden (1990:29) puts it, must be produced by the mind's own resources. Indeed, we produce representations transgressing our biologically biased views not by applying new reasoning patterns to subject matters, but by applying different but familiar reasoning patterns to the ones we are predisposed to apply to those subject matters. We can, in other words, only form new and transgressing representations by changing the mix of ingredients at our disposal, not by introducing new ingredients.

In this sense, Ruse (1986) argues that while the products of science (i.e. the representations or theories it produces) transcend their organic origin, the methods science employs and the principles it adheres to are still firmly rooted in our biology. Our scientific endeavours, in other words, as far as they can take us away from our uncritical commonsense assumptions, still flow through biologically channelled modes of thinking imposed on us by evolution (149). How else, indeed, could we think about the world, other than with the particular reasoning abilities we possess?

Ruse (1986) refers to these innate reasoning patterns underlying all human thought as epigenetic rules. This term, borrowed from Wilson (1981), designates the biological constraints on human cognition and behaviour as having their origin in evolutionary needs (Ruse 1986:143). A good example of such rules is the universal human classificatory schema into which colours are broken up (143-144), or the incest barriers we find in all human cultures (145-146). Culture, in this context, I argue in agreement with Ruse, is not some special disembodied phenomenon but 'the flesh on a biological skeleton', where the bones of that skeleton are epigenetic rules, controlled by our genetic constitution and fashioned by the hand of natural selection (147).

Those epigenetic rules include the cognitive ingredients at our disposal; the building blocks of human reasoning. According to Ruse, they are the basic logical principles we adhere to in our reasoning, such as the law of the excluded middle (either it is raining or it is not raining) and of non-contradiction (it cannot both rain and not rain at the same time), the rule of modus ponens (if it rains we stay at home; it rains so we stay at home) and alternation (either we go out or we stay at home; we don't go out so we stay at home) (Ruse, 1986:156-157). Furthermore, they include the basic premises and principles of mathematical thinking, our ability to draw causal relations, inductive and deductive reasoning, and so on. (158). Pinker (2007), on the other hand, points to our reasoning patterns about space, time, substance and causality as the substrate of human thought (cf. Chapter 3 – 4.2 Analogy and metaphor).

All human thinking, from the uncritical folk sciences we produce to the most advanced and counter-intuitive scientific conjectures, employs these basic 'cognitive building blocks'. In the same way, we cannot experience the world other than through our particular senses; we cannot think other than through those reasoning patterns evolution has provided us with. Once again delving into Boden's (1990) A.I. jargon, those cognitive building blocks can be viewed as the action-rules which are applied to the data. Transgressing theories, as pointed out above (cf. 2.4 Reconciling openness and closure) do not employ different action rules, but rather different heuristics, recombining data and action rules in novel ways. The conceptual locations yielded by those transgressing theories remain, indeed, within our conceptual space, since it contains all logically possible combinations of the data and rules we have at our disposal. All the possible cocktails, in other words, which are based on the set of ingredients we possess.

Being the products of a blind process, the cognitive building blocks at the basis of human reasoning are contingent in much the same way as the experiential realms yielded by our senses are contingent. In this regard, just as we could perceive the world differently, we could conceptualise it differently. As Ruse (1986) points out, the epigenetic rules underlying human sciences in particular and human reasoning in general are definitely adaptive and therefore the product of natural selection. Those reasoning patterns, in other words, have been selected because they boosted the chances for survival and reproduction of our hominid ancestors in their environment. The proto-human, in this light, seeing three tigers enter a cave and two leaving, and inferring that there must still be one inside, was obviously more likely to be our ancestor than the one not endowed with these basic arithmetic skills, who thinks that the cave is now empty and enters it to set up camp (162-163). However, as argued in chapter 2, the fact that natural selection has fashioned our perceptual and cognitive abilities does not entail that they provide us with a complete and accurate representation of the world.

The way we (can) represent the world, therefore, is highly dependent on the particular evolutionary road we have taken. Imagine, in this light, that we did not possess the innate faculty to represent numerical information, According to Spelke (2003) there are two distinct innate systems at the basis of our ability to represent numbers. With respect to small numbers (up to about 3) we are endowed with the innate faculty to represent the numerical identity exactly, as well as the effects of adding or subtracting one. With respect to larger sets of numbers, we represent their approximate numerical magnitudes, enabling us for instance to gauge that a set of 50 is larger than a set of 25 (but not that a set of 31 is larger than a set of 30) (297). These two distinct systems of representation are then combined by the human mind (cf. Chapter 3 – 4.1 Mapping across domains), underlying our ability to represent larger numbers exactly and therefore to count and engage in more complex mathematics (302).

Without these innate faculties underlying our sense of numeracy, therefore, not only could we never have developed mathematics, but all modern scientific representations of the world would be unreachable. Indeed, as science relies heavily on mathematics, we could never produce those theories. Furthermore, if we were unable to subdivide space and time into units, those basic frameworks through which we represent the world (we locate everything in the world in space and time) would be conceptualised in a radically different way, changing our representations at their very core. We wouldn't, for instance, be aware – at least not in a quantifiable way – of the time we lived, the time we can expect to live, the distances we cross and the time it takes to cross them.

Therefore, lacking such a cognitive ability or set of epigenetic rules, a whole series of representations would fall outside our conceptual space. We could never produce these representations simply because we lack the appropriate action-rules to be performed on the data. Similarly, we can imagine a creature endowed with more cognitive building blocks; more action-rules it can apply to its input. This would yield entire new realms of representations that are now inaccessible to human beings. A creature with such extra reasoning patterns would look upon us in much the same way as we would look upon those numeracy-lacking humans: as hopelessly closed off to some fundamental ways of representing the world. Indeed, just as there are more data in the world than can be perceived by our contingently evolved senses, there must be more ways of conceptualising this input than can be done by our contingently evolved cognitive abilities.

In this sense, the particular cognitive abilities that evolution has endowed us with limit the possible representations the human mind can produce, by providing us with a set of cognitive building blocks or innately grounded reasoning patterns that underlie all our representations. Lacking any of those building blocks would exclude a whole series of representations we now hold. Conversely, having more or different reasoning patterns would open up new realms of ways to represent the world, which are now irremediably beyond our ken. We could, as pointed out above, very well be those a-numerical beings for some other cognitive creatures with a more extensive set of 'epigenetic rules'.

Furthermore, rather than merely limiting the representations we can produce, our particular cognitive abilities also limit the possible representations we can grasp. In this sense, not only does our cognitive apparatus restrict our creative ability to represent the world, it also determines the kind of representations that are intelligible to us. Imagine we attempted to explain basic mathematics to those a-numerical mind. Wouldn't it simply be impossible for them to make sense of it? It would be an utterly hopeless endeavour, it seems: like trying to teach other primates to communicate in flawless English.

Learning, indeed, is not merely a passive process of acquiring new representations, but requires the proper cognitive structures to couch the information one is exposed to. In this regard, we might have a harder time grasping relativity theory than subtraction and addition, or the grammatical structure of our mother tongue, but we can still extend our grasp to include it by applying our cognitive building blocks. Indeed, relativity theory uses advanced mathematics, which can ultimately be broken down into reasoning patterns that we are endowed with, such as, for instance, the rules underlying logic and arithmetic (cf. epigenetic rules). A theory requiring radically different reasoning patterns, however, would simply be beyond our grasp, because we wouldn't possess the proper cognitive structure to couch these representations.

Therefore, even if we were handed representations outside of our conceptual space on a platter by some intelligent aliens, we would simply not have the tools to grasp them. We would be like those language learning chimps - unable to grasp the correct grammatical structures underlying human languages, no matter how hard our tutors try to teach us. It is, indeed, absurd – as Pinker (1994) points out – to try to have a different species emulate our instinctive form of communication. How could they succeed without the proper cognitive module to acquire human language? Similarly, it would be absurd for these aliens to try to teach us sciences for which we do not have the proper cognitive tools – as absurd as trying to teach a creature to perceive the world in a different way than its senses permit.

Indeed, as pointed out, transgressing theories overcome our commonsense representations. They do not, however, overcome our innately grounded reasoning patterns. We can, in other words, alter our beliefs about the world, but cannot alter the cognitive tools we apply to reach those beliefs. As Ruse (1986:149) puts it, while the products of our scientific enquiry may transcend their biological origin – meaning that those theories transgress our folk theories which are shaped by natural selection to provide us with a useful framework to guide our interaction with the environment (cf. Chapter 1 – 3.2.3 Species-specificity) – its methods and principles are still firmly rooted in biology. Therefore, while we are not confined to representing the world in a particular way or even set of ways, we are still restricted to representations which can be grasped through the particular set of reasoning patterns we possess, which can, in other words, be deconstructed into the particular set of cognitive building blocks we are endowed with.

4. Conclusion

Gauging the scope and limits of the representations the human mind can produce, I argue that, while we are not restricted to representing the world in a particular way or set of ways, our perceptual and cognitive nature limits the representations we can produce. We can, as pointed out above, either be biased against or straightforwardly closed off to certain representations of the world. In terms of Boden's (1990) notion of conceptual spaces, some representations are situated at the outer edge of our conceptual space, making them extremely hard to reach, while others are simply outside this space, and cannot possibly be conceptualised by the human mind.

This tension between openness and closure and transgression and limits characterises the human mind. We can represent the world in an inexhaustible amount of ways, yet these representations constitute but a particular subset of all (theoretically) possible representations. We can transgress the views we are predisposed to hold in virtue of our biology, but can only do so based on reasoning patterns which remain firmly grounded in our cognitive nature. In this regard, there is no escaping a certain biological determinism in our attempts to theorise about the world. This, however, does not mean that we have to represent the world in a certain way. Indeed, we are endowed with an open epistemic relation to the world, and are able to conceptualise it in an unlimited amount of ways.

This characterisation of the human mind as endowed with an open epistemic relation to the world, which is nevertheless comprised within a particular conceptual space, holds interesting implications for some fundamental epistemological issues. Can we represent the world accurately? Could two incommensurable representations both yield an accurate representation of the same world for the respective cognitive beings that hold them? Does a particular way of representing the world not imply an incomplete way of representing it? And finally, could the world (in principle) be represented in a universal way, or does the act of representing imply a particular perspective?

In the next chapter, I will attempt to give a comprehensive answer to those questions, turning to the main question around which this dissertation is centred: what can we know about the world, given the fact that our perceptual and cognitive faculties are the product of evolution? The answers to these questions will provide a basis for shedding new light on human knowledge - a fresh perspective with which to think about its status, which tackles the issue of realism, and looks into the threat of biological relativism that looms heavily over this dissertation.

Chapter 5: Implications for human knowledge

1. Introduction

In chapter 4, I established that while we can indefinitely mould new representations, which underlies the open cognitive relation we entertain with the world, we are nevertheless biased against and straightforwardly closed to certain theoretically possible representations of the world. In this chapter, I will look at the implications this 'open cognitive relation within the boundaries of a particular conceptual space' holds for our epistemic endeavours.

First of all, I will ask whether the fact that we possess only a particular subset of all theoretically possible ways of representing the world entails that we can only represent the world to a limited extent. Does, in other words, the fact that we have only a particular set of cognitive grasps at our disposal entail that we can only represent a limited portion of the world?

Arguing that a particular set of cognitive grasps does not entail – in principle – that our representations cannot correspond to external objects and properties, I will characterise this correspondence as a particular 'fit' determined by our cognitive nature. Outlining the consequences of our epistemic relation, where external objects are represented by means of particular cognitive tools in the light of a particular cognitive orientation, I will then ask whether this forces relativism on us.

Finally, I will look at three potential sources of limitations to a successful epistemic relation to the world – i.e. a relation where our representations fit external objects: firstly, the fallibility of our epistemic endeavours; secondly, the limited computational capacity we can bring to the task and, finally, the limited scope that could result from the particular computational tools at our disposal.

2. What is the relation between our representations and the external world?

2.1 Saving epistemological realism: the distinction between grasp and content

In the previous chapter, I established that our particular perceptual and cognitive abilities provide us with a particular and contingent set of possible representations we can hold about the world. The representations of the world, in other words, which the human mind can possibly conceive of, are but a subset of all the theoretically possible representations about that same world. This raises the question as to whether our representations actually represent the external world and – if so – to what extent. How, indeed, one could question, can a limited and contingent set of available representations represent a mind-independent reality in more than a limited way?

This question as to whether our representations of the world are about mind-independent properties and objects of the external world is at the core of the debate about *epistemological realism*. Epistemological realism – not to be confused with metaphysical or ontological realism, which is concerned with the existence of an external world – holds that our representations of the world can refer to mind-independent, external objects.[33] Epistemological anti-realism, on the other hand, holds that there is no link between our conceptual grasp of the world (i.e. our representations of the world) and the external world. In this chapter, I will argue that epistemological realism is tenable, even though we are endowed with a particular set of contingent representations of the world.

From an evolutionary perspective – claiming that our perceptual and cognitive abilities are contingent products of a blind process – it does, indeed, follow naturally that we have to admit to the possibility of radically different representations, which would be utterly unintelligible to us. In this context, several philosophers have argued from evolutionary considerations that our grasp of the world is but a particular one. This point is often framed as the possibility of 'alien scientists'. Clark (1986) points at the 'interesting consequence' that

[33] Epistemological realism is defined by the Theological and Philosophical dictionary as entailing that 'the mind knows independent things not ideas alone'. Epistemological realism, in other words, entails that our representations refer to external, independent objects. (Cf. http://www.philosophy-dictionary.org/Epistemological_realism).

we must accept the possibility of alien epistemologists, working successfully with a different model of our 'common reality'. Indeed, he argues, 'the ideal limit of human scientific enquiry is still not the only possible 'correct' representation of reality even if relative to our cognitive constraints and observational access there are no visible alternatives' (158).

Similarly, Rescher (1990) argues that 'there is no categorical assurance that [alien] intelligent creatures will think alike in a common world, any more than they will act alike – that is, there is no reason why *cognitive* adaptations should be any more uniform than *behavioral* adaptations' (his italics, 92). Sciences, he continues, 'are bound to vary with the cognitive instruments available in the physical constitution and mental equipment of their developers' (95). Our sciences, in this regard, are but the intellectual product of one particular sort of cognitive life-form. They are ultimately species-relative (95).

The most important reason that we cannot expect alien scientists to engage in our kind of science is, according to Rescher, that the possible sorts of natural science are almost endlessly diverse. Natural science – in the sense of 'an inquiry into the ways of nature' – is, Rescher argues, 'endlessly plastic' (94). Its development is bound to reflect 'a historical course closely geared to the specific capacities, interest, environment, and opportunities of the creatures that develop it' (94). It is, in this regard, not a process that must follow a similar route to ours and end up with a similar product. In this context, Rescher concludes that 'it would be grossly unimaginative to think that either the journey or the destination must be the same – or even substantially similar' and that 'unless we narrow our intellectual horizons in a parochially anthropomorphic way, we must be prepared to recognize the great likelihood that the 'science' and 'technology' of a remote civilization would be something very different from science and technology as we know it' (94).

This crucial implication of bringing the theory of evolution to bear on the question of the limits and scope of human knowledge could lead us to think that we have no or only limited cognitive access to the external world. The main argument against epistemological realism is, indeed, what I'll refer to as the *conceptual screen argument*. The conceptual screen argument starts from the claim that we can only grasp the world through our own particular conceptualisations. I fully agree with this claim – our grasp of the world is indeed the product of the particular and contingent set of conceptual resources we have at our disposal. However, from this premise, the conceptual screen argument derives that it is impossible to 'see' the world as it really is. In other words, because we must grasp the world by means of a particular set of conceptual resources, we can never grasp it as it really is. Epistemological realism, therefore – this position claims – is untenable. This conclusion, I reject.

Indeed, the conclusion drawn by the conceptual screen argument implies that true knowledge of the external, mind-independent world would require neutral access to the world, untainted by the subject's perceptual and cognitive 'spectacles'. Whenever we view the world through our own conceptual frameworks, it argues, we do not view the world-in-itself but merely the Kantian phenomenal world - the world as it appears to us. This premise, which I will be challenging, is at the basis of the traditional epistemological project of *foundationalism* - the attempt to ground all knowledge in basic, un-doubtable (i.e. untainted by the cogniser's perceptual and conceptual framework) statements about the world.

Foundationalism, however, is doomed to fail. The very possibility of direct epistemic access (i.e. neutral, unmediated or objective access) is problematic. Indeed, either we take this unmediated 'given' as arising from perception – as the empiricists did – or we take the rationalist approach and consider cognition as the origin for this 'given'. From an evolutionary perspective, however, both are untenable. As pointed out in chapter 1, perception does not yield a subject- or species-independent picture of the world – we perceive but a small part of the world in a particular way. Cognition, on the other hand, yields a set of possible ways of representing the world that are as equally contingent and particular as our perceptual abilities (cf. Chapter 4). There is, in other words, no way to gain epistemic access to the world, other than through the particular 'lens' evolution has provided us with.

Furthermore, as Sellars (1956) objects, the very idea that we can ground knowledge in cognitive states, which are in direct contact with reality, is problematic. This foundationalist picture, he argues, requires, on the one hand, that there must be basic cognitive states, in the sense that those states are independent of any epistemic relations to other cognitive states. On the other hand, it requires that all non-basic cognitive states are derived from those 'given' basic states. Sellars's argument is not aimed at denying that there can be basic cognitive states, but in denying the possibility of a foundationalist structure of knowledge, because – he argues – the standard candidates for basic states (e.g. sense-data) either fail to be independent from other cognitive states (because they presuppose other knowledge) or, if they truly are independent, fail to support other cognitive states.

Indeed, to satisfy the first condition, which Sellars labels the 'Epistemic Independence Requirement', the basic cognitive state must be independent of inferential connections to other states. To satisfy the second condition – the 'Epistemic Efficacy Requirement' – basic cognitive states must, on the contrary, participate in inferential relations with other states. They must, therefore, possess propositional form and truth-value. According to Sellars (1956), however, 'examination of multiple candidates for non-inferentially acquired, propositionally structured cognitive states indicates that their epistemic status presupposes the possession by the knowing subject of other empirical knowledge, both of particular and of general empirical truths'[34] (cf. de Vries, 2011). Since it presupposes the possession of other empirical knowledge, it cannot be a basic cognitive state.

According to de Vries (2011), since 'there is no exhaustive list of all possible candidates for the given, the argument is not a conclusive, once-and-for-all refutation of foundationalism, but it is a significant challenge to it, putting the burden of proof on the defenders of a given'. The very enterprise of foundationalism is, therefore, doomed to fail, not only because of the impossibility of having direct epistemic access to the world, but also because – even if this were possible – it is, according to Sellars, very unlikely that we will ever come up with cognitive states which are both basic *and* efficacious.

Does this mean we have to yield to the conceptual screen argument and give in to anti-realism, accepting that we have no epistemic access to the external world? I argue – as signposted earlier – that we do not. The fact that we represent the world through a particular, species-specific framework does not entail that those representations do not represent elements of the external world. Claiming that it does is the result of what Nagel (1986:101) calls a confusion between form and content. As Nagel argues, we need to distinguish between

[34] Giving a full overview of Sellars's examination of proposed basic, propositional cognitive states, however, is beyond the scope of this chapter.

the content and the form of a thought. The content, in this regard, is quite independent from the particular form in which it is expressed; independent, for instance, of the particular language in which it is formulated. All thoughts, Nagel continues, must have a form which makes them accessible from a human perspective. This, however, does not entail that *they are about* our point of view. 'What they are about depends not on their subjective form, but on what has to be referred to in any explanation of what makes them come true' (101-102). Their content, in other words, is determined by *what* they refer to in the world, not by *how* they refer.

The content of a grasp (I will refer to Nagel's (1986) 'form' as grasp), therefore, is not necessarily particular or species-specific because the grasp itself is. Different (subjective) grasps can, indeed, grasp the same (external) thing – i.e. have the same content. Take, for instance, two different creatures, each with a particular ability to represent a point in space. The first represents a point in space by using numerically expressed 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 *grasp* of spatial locations at its disposition 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.

'Grasp', in this regard, refers to the particular way in which we access the world. The fact that these grasps are both particular and contingent – i.e. that there are other theoretically possible ways of representing the world and that the world does not need to be represented in this way in order for the representation to be accurate – does not, as I argued in the previous chapter, entail that the content of that grasp is 'deformed' by the conceptual apparatus of the cognitive creature. Indeed, as pointed out in the example above, two different or even mutually unintelligible grasps can have the same content. The same point can, indeed, be represented on a numerical Cartesian axis system and in a grid system.

In this regard, the fact that we have a particular grasp – i.e. a particular set of representational tools – at our disposal does not imply that this grasp cannot access elements in the external world. Nor does the possibility of unintelligible representations of the world – to the human mind, that is – which follows, as pointed out, from an evolutionary take on our perceptual and cognitive faculties (cf. the hypothesis of the possibility of alien scientists), entail that our representations, at best, represent the world in a limited way. It is not because we look at the world from a particular vantage point (i.e. that other vantage points are theoretically possible) that we can't have a clear view.

2.2 The nature of cognitive grasps: impossibility of a God's eye view

Thus far, I have established that although we have a particular set of grasps at our disposition, this does not entail we have no epistemic access to the world. However, we can take the argument further by reflecting on the nature of cognitive grasps. In this regard, rather than merely claiming that we grasp the world in a particular way, I will argue that every (theoretically) possible cognitive grasp (even a 'perfect' one) takes a particular perspective. In other words, the particularity of our perspective on the world is not merely a product of our (limited) cognitive nature, but follows necessarily from the very act of representing. 'Representing', in this regard, refers to the cognitive ability to 'translate' something through a particular medium. We can represent a tree, for instance, through a drawing, a statue or even a linguistic description. In doing so, we translate the features of that tree either to a two-dimensional colour arrangement (a drawing), a three-dimensional matter arrangement (a statue), or a set of conventional symbols (a linguistic description).

A cognitive grasp is a representation. It 'translates' external elements into an internal, mental medium. In doing so, it re-presents them – i.e. makes them present again – in some medium. Any possible kind of cognitive grasp which provides epistemic access to the world requires such a translation. This follows from the fact that something external is represented internally. Features and properties of the world, whether they yield sensory impressions, linguistic descriptions or mathematical equations are, in this regard, not objectively or neutrally represented – i.e. without the mediation of any structuring framework of the cogniser – but are 'translated' in a particular way. There can, in this sense, be no such thing as a universal or objective grasp, since the very nature of a cognitive grasp entails that it provides a (particular) medium in which properties of external objects are represented.

Furthermore, as argued above, there is no one-to-one correlation between a grasp and its object. Every grasp discloses a particular content (i.e. represents a particular external object or property) but external objects and properties do not require a singular grasp to be represented. They can be grasped in more than one way. Indeed, as the example shows, the same point in space can be represented or grasped either by means of a numerical Cartesian axis system or by means of a grid system. In this context, I would like to cast Quine's (1951, 1975) underdetemination thesis in a broader perspective.

Underdetermination is formulated by Quine (1975:327) as the possibility that our postulated complete global theory of the world could have empirically equivalent alternatives that are incompatible with ours, or in Quine's words: 'with no translation from one to the other possible'. This possibility follows, indeed, from the fact that while a hypothesis (theory or representation) entails the existence of certain observable facts, those facts do not imply the particular hypothesis. As Quine (1975) phrases it:

> The hypotheses are related to observation only by a kind of one-way implication; namely, the events we observe are what a belief in the hypotheses would have led us to expect. These observable consequences of the hypotheses do not, conversely, imply the hypotheses. Surely there are alternative hypothetical substructures that would surface in the same observable ways (313).

While Quinean underdetermination holds that the world does not determine a single theory – i.e. that even within the framework of a singular set of grasps (human cognitive grasps, for instance) different theoretical constructs could still account for the data – I argue, given the aforementioned possibility of alien scientists, that the world does not determine a unique set of grasps. Since, as pointed out above, the same objects and properties of the world can be grasped in various ways, any possible grasp would have to allow for the existence of different, equally 'fitting' grasps. In this regard, even a perfect theory of the world – a theory that explains the cosmos both in its totality and in its every detail – would still have rival cognitive grasps accounting for at least parts of the data it explains. Indeed, since there is a one-way implication between grasp and object, states of the world do not determine singular representations corresponding to them, but allow for a variety of different cognitive constructs to be represented.

In this regard, the traditional epistemological ideal of viewing the world from nowhere – or from everywhere for that matter, if we take a God's eye view to comprise all possible (accurate) perspectives – is misguided. Epistemological realism (claiming that representations represent external objects accurately) does not require 'an impossible attempt to view the world from nowhere' (cf. Putnam, 1990:28), and nor does an accurate grasp of external reality entail that we shed the lenses through which we view the world. This neo-Kantian epistemic pessimism is the result of what Kaufman (2002:343) calls 'the opaque screen view of conceptual schemes' – the view that the conceptual schemes we must employ shut us off from the true nature of the world.

This view is at the core of epistemological anti-realism, or 'quietism', as Wright (1992) puts it, which claims that epistemic access to the world is prevented by our inability to look beyond our conceptual schemes. This, according to thinkers like Putnam (1990) and Rorty (1979), makes any claim about the external world impossible. There is, according to these philosophers, no sense in talk that looks beyond our conceptual frameworks (Kaufman, 2002:343). I propose a diametrically opposed position. Rather than shutting us off from the world, the conceptual schemes or the particular grasps – to use the terminology defined above – we have at our disposal are a necessary condition to gain epistemic access to the world.

Indeed, as pointed out above, the nature of a cognitive grasp – i.e. the fact that it represents something external in a mental medium – implies that it is a particular grasp. Every possible way of grasping the world would involve conceptual schemes or frameworks of some sort, since they provide the medium in which external elements can be internalised. Without such particular ways of representing the world, there simply cannot be any kind of cognitive relation between an organism and its environment. Therefore, rather than being opaque blindfolds shutting us off from the world, our particular cognitive grasps are actually windows onto an otherwise invisible world. Neo-Kantian epistemic pessimism, ranging from epistemological anti-realism (there in an outside world, but we have no epistemic access to it) to outright metaphysical anti-realism or idealism (there is nothing outside our inner experience) is, in this regard, caused by a denial of the true nature of the act of representing. The particularity of the perspective through which we represent the world is not an obstacle, but, on the contrary constitutes the very nature of a (or any, for that matter) cognitive relation to the world.

2.3 The fallacy of correspondence as 'mirroring'

This brings to light the core fallacy of the traditional epistemological view of knowledge. Knowledge, it is argued, should be an objective *reflection* of its object - it should *mirror* the external world. The mirror metaphor of true representations entails 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 (1922). 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.

Nevertheless, as Goldman (1986) points out, this invites the objection that 'the world does not contain factlike entities'. It does not contain entities which would correspond exactly to sentences or propositions. Language and thought, Goldman argues, construct disjunctive, existential, conditional, and universal statements, and so on. Supposing, however, that the world contains disjunctive, conditional, existential and universal structures, is misguided. The Tractarian version of the correspondence theory – portraying the world as being structured in truth-like entities – is therefore, according to Goldman, an untenable doctrine (151).[35]

[35] Whether Wittgentein committed to this view in the Tractatus is not entirely clear. However, I will follow Goldman (1986) by referring to the view of correspondence as mirroring as the 'Tractarian version of correspondence'.

I fully agree with this analysis, based on the aforementioned point that we cannot confuse grasp and content. Arguing that our propositions mirror external facts and properties amounts to reducing the world to the particular grasp we have of it. Thought and language are the cognitive tools enabling us to grasp and represent elements of the external world. Deriving from this that the world must be structured identically to human thought and language is the wrong way of looking at things. It is like arguing that trees are ultimately colour arrangements, or that space must be 'gridded', because one is accustomed to represent trees in a drawing or points in space in a grid system. The world does not reflect the properties of our grasp, but our grasp constitutes a medium in which properties of the world can be represented. Confusing this one-directional relation between external objects and representations (representations express properties of the world, but the world does not harbour the structures of the representational medium) for a symmetrical or two-way relation is fundamentally mistaken.

This problematic character of this criterion of correspondence – claiming that true propositions (must) mirror their objects – is often invoked as a persuasive argument against the possibility of any kind of correspondence between representations and the external world. Nevertheless, as Goldman (1986), argues, the mirror metaphor is only one possible metaphor for correspondence (152). While this particular criterion for correspondence is untenable, there is no reason to abandon correspondence altogether and sever all ties between internal representations and external objects.

In this regard, Goldman (1986) 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, therefore, that the world can be interpreted in an infinite number of ways, as Goldman argues, does not entail that a particular way cannot be right or wrong. There is, indeed, a connection between representation and its object: the representation is or should be 'moulded' to fit the 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.

Therefore, the fact that we are bound to a particular and contingent conceptual space does not entail that we cannot – or can only to a limited extent – represent the world. Our representations can *fit* their objects, just as both a drawing and a statue can depict an object. The claim that there can be no correspondence between representation and its external object rests on the assumption that the utterly untenable mirror metaphor is the only possible criterion for correspondence. Goldman's (1986) metaphor of correspondence as fittingness, however, enables us to ground epistemological realism by acknowledging that we have epistemic access to the external world, while admitting that the human cogniser must represent the world in a particular and contingent set of ways. One question, however, remains unanswered. What is the nature of this *fit*?

2.4 The nature of the fit

In chapter 3 I pointed out that we are endowed with an epistemic value system (cf. Chapter 3, 6.3 Epistemic value system). This comprises both an epistemic goal and epistemic criteria for realising this goal. Such a goal, I argued, is our predisposition to look for truth. Indeed, as Papineau (2000) claims, our search for truth is an innate drive, characteristic of our species' curiosity, our hunger for knowledge and our need for justification before adopting a belief. Truth or justification, however, as I pointed out, presuppose some criteria in the light of which they can be realised. Such criteria are the epistemic values we adhere to: predictive accuracy, coherence, scope and simplicity. Our epistemic value system or epistemic orientation, in this regard, endows us with a set of directives; namely, those epistemic values that guide our epistemic endeavours, drawing us towards certain representations while rejecting others.

Nevertheless, when representing the world, we often transgress some of those values. In this regard, we might, for instance, ignore the value of simplicity when it appears that those neat and simple hypotheses do not produce accurate predictions or lead to inconsistencies within the theory or – vice versa – we might ignore predictive accuracy to accommodate simplicity and scope. Chalmers's (1976) account of the Copernican revolution is interesting in this regard. Still immersed in Aristotelian mechanical and cosmological thinking, Copernicus's contemporaries came up with some strong arguments against the Copernican theory of a moving earth. The main counter-argument is the so-called 'tower argument'. If the earth spins on its axis, Copernicus' contemporaries argued, a stone dropped from the top of a tower (erected on this spinning globe), should strike the ground at some distance from the tower, since the tower – sharing the spinning motion of the earth – will have moved by the time the stone reaches the ground. Copernicus himself, as Chalmers points out, had no adequate response to those arguments, which seemed to refute his theory. Nevertheless, his view was not simply dismissed, precisely because it offers such a neat way of explaining planetary motion. Whereas the Ptolemaic account of the universe needed to ascribe a whole series of complex motions to planets to account for the data we gather from them in observation (such as, for instance, planets retracing their path in opposite directions), Copernicus's view could explain the observational data in a much simpler and neater way (95-96). Therefore, although the Copernican system failed in respect of predictive accuracy (the stone did not fall at some distance from the tower), it was however not dismissed, because of the simplicity it boasted in describing the planets' motion.

In this regard, we transgress epistemic criteria in the light of other criteria. We ignore simplicity to accommodate coherence or we ignore coherence to accommodate predictive accuracy, but we cannot do away with all these values at once. This would – indeed – leave us utterly blind in our conceptual space, with no possible way of asserting that, for instance, the big bang is a better theory for the origin of the universe than the Cherokee beetle diving into the endless waters and pushing up some mud.

This brings Neurath's (1959:201) metaphor to mind of having to rebuild a ship while staying afloat. Arguing that we cannot reason from a tabula rasa, he famously exclaimed that 'we are like sailors who must rebuild their ship on the open sea, never able to dismantle it in dry-dock and to reconstruct it there out of the best materials'. Indeed, we can only criticise and overcome parts of our knowledge about the world based on other knowledge. In this sense, when we transgress certain epistemic values we can only do so supported by the remaining bulk of those values. Dismantling them all at once would sweep the ground underneath our feet or – to endorse Neurath's nautical metaphor – would take out the lower deck and sink the ship.

Furthermore, when partially transgressing our epistemic values, we do sacrifice the cognitive satisfaction that such a representation provides us with. Our mind, in other words, is biased against representations which do not fulfil all epistemic values we find ourselves endowed with. In this context, a lack of coherence, or the resistance to a simple and elegant solution, is often perceived as an obstacle to be removed in order to achieve a comprehensive theory. A good illustration of such an epistemic value transgressing representation is the wave-particle duality postulated in quantum mechanics. According to the Copenhagen interpretation, the behaviour of quantum scale objects cannot be reduced to either particle-like or wave-like behaviour. The phenomena can be viewed in one way or another, but not in both simultaneously. In this case, the value of coherence is sacrificed to accommodate the data (i.e. to accommodate the value of predictive accuracy). This, however, makes the theory imperfect for us, inviting widespread criticism since it is hard to accept that it could provide us with an entirely truthful and accurate representation of its subject matter.

These values are, indeed, part of our innate endowment (cf. chapter 3, 6.3 Epistemic value system). As Carruthers (2006:347) points out, they are universal to all human cultures from hunter-gatherer societies to Western scientific communities. In this regard, the criteria or values we adhere to for realising our epistemic goal, are firmly grounded in our cognitive nature. They are – like any other genetically determined feature – the particular and contingent outcome of a blind evolutionary process.

This brings us to the nature of the fit. We perceive representations as accurate *in the light of* those epistemic values. The way we fit representations to their (external) objects, therefore, is by means of a particular, innate set of values or criteria fulfilling a cognitive goal – i.e. truth or accuracy. Therefore, not only are we endowed with a particular and contingent set of representational tools that we use to produce fitting representations of the external world, but the fit itself is also contingent and particular. In other words, when representing the world, we bring both a particular set of tools to the task *and* a particular set of rules by which the (external) object of representation is translated into the representation. Just as a painting (i.e. the representation) of a tree (i.e. the object of representation) requires both the use of a set of tools (paint, white canvas) and a set of rules (perspective, proportions, colour arrangement, etc.).

Rescher (1990), in this context, argues that we are endowed with a particular 'cognitive project' which is 'the intellectual product characteristic of one particular sort of cognitive life-form' (95). Indeed, not only the available cognitive tools depend on the biological endowment of their developers but also the purpose to which these tools are put to use. The set of epistemic values determining which representations fit the world for the human cogniser, in this regard, are as much a given as its perceptual and cognitive abilities. Both cognitive purpose and cognitive means are ultimately species-specific.

From an evolutionary perspective, therefore, not only do we have to envisage the possibility of (radically) different ways of representing external elements (i.e. other representational tools) but also the possibility of different fits between representation and object. Our alien scientists could, therefore, represent the world in a radically different way, both because the 'fabric' of their representations is different *and* because the criteria in virtue of which representations are desirable – i.e. are perceived to 'fit' their objects – are different. The same object X could be represented by a cognitive being A by means of representational tools C,D,E in the light of criteria F,G,H and by cognitive being B by means of I,J,K in the light of L,M,N. One cognitive being could, for instance, represent gravity by means of a mathematical equation in the light of the criterion of predictive accuracy (in what time an object of a certain mass will drop from a certain height) and another could represent it as a non-quantified latent force proportional to mass, in the light of the criterion of explanatory scope (accounting for the fact that objects are subject to the earth's gravitation, the earth to the sun, our galaxy to other galaxies, and so on).

In this context, with regard to the earlier statement that two different cognitive beings can represent the same object *accurately* in a radically different way, it is important to note that what makes both representations accurate – or rather 'fit' – is not some objective given, but is equally grounded in their respective cognitive make-up. A particular representation, therefore, can be said to correspond (i.e. fit) to an external element or property of the world *for* a certain cogniser, if the representation aligns with the epistemic orientation (i.e. the particular criteria determining the fit) of the cogniser, given the characteristics of the external object which is represented. To sum up, the epistemic correspondence between our representations and the world is a particular fit with particular tools.

3. What are the consequences of our epistemic relation?

3.1 The two-sided determination of knowledge

To recapitulate, our (true) representations of the world are not objective reflections of the external world, but particular fittings – in the light of our epistemic orientation (i.e. the set of epistemic values or criteria we are endowed with) – with available representational tools. These representations do, therefore, not represent the world in vacuo, but only in reference to the particular kind(s) of available representations and the epistemic orientation, both grounded in our cognitive nature. Our representation of the world, in this regard, is inexorably linked to our biologically determined cognitive nature. It is a representation of external objects and properties in the light of a given orientation by means of a given set of tools.

Since what constitutes knowledge of the world for us is determined both by states of the world (it must fit the constraints of those external objects and properties) *and* by the nature of our cognising apparatus, it makes no sense to talk about knowledge independently of the type of cognisers we are, for whom a particular set of available representations represent an external object (accurately) in virtue of a particular epistemic orientation. Our knowledge of the world, in this regard, is not some outer-worldly, universal body of representations held by any creature that would be endowed with an inkling of thought, but a species-specific product determined by our cognitive apparatus.

This makes sense in an evolutionary light, where our perspective on the world can be characterised as a view from within rather than a view from outside. As Munz (1993) phrases it: 'in biological perspective, the knower is no longer an observer from outer space, looking upon a world in which there is nothing but swirling electrons or sub-atomic particles. The knower, is, on the contrary, an integral part of the known system' (191). Looking at the world from within, we do not capture it in some neutral, objective or direct way, but in the particular way determined by our position within the 'system'.

To put it differently, evolution stumbled upon cognitive adaptations, providing organisms with the ability to internalise external properties of the world, in the same way that evolution stumbled upon visual adaptation – i.e. eyes – enabling organisms to (among other things) 'navigate' through space. Just as eyesight is not the only possible way to navigate through space – consider a bat's echolocation – our cognitive apparatus does not provide us with the only possible way of understanding the world. Other cognitive apparatuses could entertain a radically different epistemic relation to the same world, or, to stay with the metaphor, other positions within the system are possible (cf. hypothesis of alien scientists). Furthermore, this entails – and this is the point I want to make – that the products of any given cognitive system only make sense in relation to the cognitive system that produced them, just as a sound bouncing off walls only constitutes navigational information for bats, and photon particles bouncing off surfaces only constitute navigational information for organisms relying on eye-sight.

3.2 The threat of relativism

This raises the threat of relativism. Since our knowledge of the world is inexorably linked to our cognitive nature, it can be said to be relative to that cognitive nature. What constitutes knowledge for us is – indeed – the product of our epistemic orientation. A representation, therefore, is true or false in virtue of our biologically determined epistemic orientation (i.e. our epistemic values or criteria), not in virtue of a framework that transcends our contingent cognitive nature.

The fact, however, that the nature of the fit (between our representations and the world) is determined by our innate epistemic orientation, can hardly be seen as undermining the possibility of knowledge. Indeed, it would be misguided to consider our particular and contingent cognitive nature as a threat to the possibility of acquiring knowledge of the world, since the criteria determining whether a representation fits the world, and therefore constitutes knowledge for us, are rooted in our cognitive make-up. Knowledge, in this regard, is itself a contingent product of our evolved minds. It is defined in terms of the criteria grounded in our cognitive apparatus. Stating, therefore, that true knowledge of the world is prevented or limited by the 'distorting perspective on the world' that our cognitive nature provides us with makes no sense, since what constitutes knowledge for us is determined by our very cognitive nature.

Making an even stronger point, I will argue that, rather than raising the threat of relativism, the biological determination of knowledge actually wards off relativism. Indeed, the fact that knowledge is rooted in our cognitive nature provides us with the necessary grounding to rebut the traditional forms of cultural relativism. Theories positing a form of epistemic relativism, stating that all knowledge of the world is determined by or relative to some framework, indeed typically take this framework to be a cultural product. The way we view the world, relativists claim, is the product of the particular beliefs of the culture we are raised in, the particular language we possess, or the particular scientific paradigm we adhere to.

Cultural relativism can either take a descriptive or normative form. While descriptive relativism consists in the empirical claim that different cultures, linguistic or scientific communities have different core beliefs, modes of thought or standards of reasoning, normative relativism claims that core beliefs, modes of thought and standards of reasoning are *only* right or wrong in the light of those cultural factors. Normative relativism states, in other words, that there is no other (higher) framework determining whether our representations of the world are accurate, other than the particular cultural frameworks we possess (Swoyer, 2003). For the purposes of this chapter, I am only interested in this stronger, normative version.

Normative cultural relativism can take this determining framework to be a particular set of background beliefs, the particular language, or the scientific paradigm in which the cogniser is raised. In this regard, Franz Boas (1887:589) exclaimed that 'our ideas and conceptions are true only insofar as our civilization goes'. There's no epistemic measure, in other words, of the way we represent the world, outside the contingent framework(s) set by a particular culture. Therefore, knowledge itself becomes an utterly contingent product of a particular culture.

One of the historically most influential theories defending a cultural form of relativism is – without any doubt – linguistic relativism. This view holds that the structures (vocabulary and syntax) of a language determine – or at least affect to an important extent – the way speakers conceptualise the world. While the theory is commonly referred to as the 'Sapir-Whorf hypothesis', it can be retraced to Wilhelm von Humboldt (1820), who declared that 'the diversity of languages is not a diversity of signs and sounds but a diversity of views of the world' (Trabant, 2000).

As pointed out, the most prominent defenders of the thesis that human thought is determined by the language of the cogniser are Sapir and Whorf. Whorf, Sapir's pupil, was the first to base the hypothesis that language determines thought on empirical findings. Studying the native American Hopi language, Whorf (1956) argued that – among other differences – the Hopi language lacks any nouns referring to units of time (e.g. 'hours', 'days', 'months' and 'years') and that, consequentially, the Hopi did not and could not treat the flow of time as a sequence of different countable units. This is, of course, in stark contrast to English speakers and speakers of related 'standard European languages' – as Whorf referred to them – who, possessing nouns referring to units of time, view time as a sequence of countable units. This brought Whorf to conclude that:

> We dissect nature along lines laid down by our native language. The categories and types that we isolate from the world of phenomena we do not find there because they stare every observer in the face; on the contrary, the world is presented in a kaleidoscope flux of impressions which has to be organized by our minds – and this means largely by the linguistic systems of our minds. We cut nature up, organize it into concepts, and ascribe significances as we do, largely because we are parties to an agreement to organize it in this way – an agreement that holds throughout our speech community and is codified in the patterns of our language. [...] All observers are not led by the same physical evidence to the same picture of the universe, unless their linguistic backgrounds are similar, or can in some way be calibrated (Whorf, 1956:212-214).[36]

While strong linguistic relativism and other strong forms of cultural relativism have been largely abandoned, it remains widely accepted that any theory or view of the world is – to an important extent – the product of a given cultural perspective. This entails that there is no ultimate measure against which one view can be considered as better than another, since both are the product of contingent factors such as historical background, language or scientific paradigm.

Against these forms of relativism, I defend the view – at the basis of this dissertation – that humankind is endowed with an innately grounded cognitive nature. We possess, as pointed out, a set of cognitive tools that we apply in representing the world, as well as an epistemic orientation by which we fit our representations to the input we gather from the world. All

[36] The Sapir-Whorf hypothesis, however, came under heavy criticism. Ekkehart Malotki's (1983) study of the Hopi time expressions presented a number of compelling findings challenging Whorf's characterisation of Hopi language and culture as being timeless. Furthermore, the Chomskyan universalist theory of language (Chomsky: 1965) pointed at a 'universal grammar', shared by all natural languages. Languages, therefore, rather than being particular cultural artefacts that radically differ from one another, came to be viewed as sharing the same underlying structure. The perceived differences between languages, in this regard, are argued to be merely surface differences which do not affect cognitive processes, which are universal to humankind.

human cognisers, in this context, have a common epistemic project which they strive to realise with a common set of cognitive means. Theories about the world, therefore, are not fitted to some utterly contingent spatio-temporal framework – i.e. the framework determined by the particular culture in which they are articulated – but are ultimately fitted to a framework which is grounded in our nature. There are, in other words, not a multitude of radically different possible perspectives on the world that humankind can adopt with equal cognitive satisfaction, but a biologically determined one that it is bound to cast on the world in virtue of its nature. [37]

This, of course, does not imply that culture is not a factor in the way we come to view the world. Different cultures have radically different beliefs about the world and different areas of interest they pass on to the next generation. We obviously do not come into a world which is laid out in front of us as a virgin terrain ready to be explored without any inherited preconceptions and analyses. In this regard, it is trivially true that a Pygmee living in the central African jungle in the 19th century looks at the world differently, and explains the same phenomena in a radically different way, to a European scholar at the beginning of the 21st century.

My argument – in this regard – is not that our shared cognitive nature provides for a unique and universal view of the world shared by all humankind. While this is true for the commonsense theories we hold about the world (cf. Chapter 1 – 3. Intuitive theories), it obviously doesn't apply to our scientific or other explicitly formulated theories, which transgress these commonsense assumptions about the world. Indeed, the very fact that we are able to transgress our biologically determined views is what opened up the kaleidoscopic realm of different cultural perspectives, as argued in chapter 3. We are, as pointed out previously, endowed with an open cognitive system (cf. Chapter 4 - 2.3 Open cognitive relation), enabling us to shape an infinite number of different representations of the world.

[37] This does not, however, necessarily entail that the state of the world and the nature of our cognitive apparatus would determine a single, ideal theory or set of theories. We do, indeed, have to admit to the possibility of Quinean underdetermination (1975), stating that even our best scientific theories might be underdetermined by empirical data. In this regard, I merely commit to the view that our cognitive nature seriously constrains our 'fitting' (and therefore desirable) theories of the world. The framework in virtue of which theories of the world are right or wrong, in other words, is by no means (exclusively) culturally determined.

The distinctiveness of human cognition, in this regard, is precisely that we are not forced to look at the world in a particular way.

This being said, we do possess a unique epistemic framework rooted in our cognitive nature. In this sense, while we are not born into a virgin world, neither are we born as virgins into the world. All human beings come equipped with similar perceptual and cognitive tools that they apply to the world. Furthermore, we share a common epistemic orientation (cf. Chapter 3 - 6.3 Epistemic value system). What makes a representation true for us is – therefore – not some contingent cultural framework, but a necessary one, since this framework is innately determined.

The particular natural language we happen to acquire, in this regard, is not the inescapable prison house of thought it was once thought to be. Indeed, it is both the case that natural languages themselves are constrained by shared cognitive structures (cf. Chomsky's universal grammar), and therefore are not as radically different as previously thought, and that human thought, in general, is constrained by a set of shared abilities and a shared orientation. Similarly, modern sciences are not merely the product of some contingent cultural framework, but are constrained by universal principles (universal among human cognisers, that is). Recall, in this regard, the fact that the epistemic values modern scientists adhere to are shared by all humans, including illiterate hunter-gatherer societies (cf. Chapter 3 - 6.3 Epistemic value system). As Ruse (1986) so elegantly phrases it, our scientific endeavours 'flow through biologically channelled modes of thinking imposed on us by evolution' (149). What makes a theory right or wrong, better or worse, or fitting or not, therefore, is ultimately determined by our cognitive nature, not by the particular culture in which the theory is formulated. The framework by which our view of the world is fitted, in this regard, is written by the hand of nature, not by the hand of humankind.

Knowledge, therefore, is the product of the biologically determined cognitive relation between the particular cognisers that we are and the world. This biological determination, rather than imposing an undermining form of relativism on our epistemic endeavours, provides us with new grounding from which we can resist the destructive threats of cultural relativism. Indeed, as pointed out, knowledge presupposes an epistemic goal and a set of values which are rooted in our cognitive nature and – therefore – only makes sense within the cognitive relation between the type of cognisers we are and the world. Biologically grounding human knowledge, in this regard, takes the necessity of an epistemic framework at face value, without succumbing to an 'anything-goes' relativism. A theory fits, not in the light of a convention or a contingent cultural framework, but in virtue of our very cognitive nature.

The fact, however, that we have epistemic access to the external world, i.e. that we can successfully represent external objects and properties within the framework of our epistemic relation to the world, does not imply that there are no limits to this ability. Indeed, while our representations can correspond (as defined above) to external properties and objects, it does not mean that they automatically do so, nor does it mean that they represent the world in a complete fashion, representing everything there is to represent in every detail.

4. What are the limits to our epistemic relation to the world?

4.1 Limits set by fallibility

Epistemic fallibility, i.e. the view that human beings *could* be wrong about their beliefs and understanding of the world, is – as Artigas (1992) points out – widespread in contemporary epistemology. Indeed, who in his right mind would deny that human beings can be mistaken? While this claim is trivial with regards to our everyday experience – who doesn't forget, misinterpret or deform information? – it also applies to our more carefully crafted and deeply reflected representations of the world.

Evolution does, indeed, provide good reasons to support the view that our cognitive apparatus can lead us to misrepresent aspects of our environment. It is, as argued in chapter 2, shaped to provide us with biological fitness-enhancing behaviour, not disinterested accurate representations. In this regard, studies in cognitive science have brought to light a myriad of cognitive biases that human thinking is subject to (cf. chapter 2, 3.2.4 Cognitive biases and error management theory). Those 'heuristics', or mental strategies for information processing, can in some cases lead to erroneous inferences. While we can often detect cases of fallacious reasoning and 'straighten out' our reasoning, we cannot exclude the possibility that fallacious reasoning seeps into even our best scientific endeavours.

Furthermore, rather than merely being the product of fitness-enhancing heuristics, the fallibility of our cognitive apparatus follows more generally from a natural perspective on the human mind. Assuming that we are not bestowed with some outer-worldly, all-potent rationality, but that we have a set of evolved perceptual and cognitive abilities that we apply to represent the world, we have to come to terms with the fact that – while they can and do provide us with accurate (i.e. fitting) representations – they can also err and fail to represent. Just like any other feature of any other organism, our cognitive abilities are imperfect and limited. In this regard, McGinn (1989:350) claims that: 'Representational power is not all or nothing. Minds are biological products like bodies, and like bodies they come in different shapes and sizes, more or less capacious, more or less suited to certain cognitive tasks'.

The fallibility of the human mind is, in this context, confirmed by the history of science, revealing a number of instances in which our representations of the world did not fit the actual states of the world. For example, it was once commonly accepted that the earth was flat and at the centre of the universe, that burning substances release 'phlogiston' – an element without taste, mass, odour and colour – and many scientists right up to the 19^{th} century believed in the possibility of 'spontaneous generation', claiming that life could arise from inanimate matter (supported, for instance, by the spontaneous appearance of maggots from rotting meat). All these theories, as Laudan (1981) points out, were both successful in the past – i.e. they seemed to be empirically confirmed – but were non-referential – they did not refer to an element in or property of the external world (33). The world, we discovered, is neither flat, nor in the centre of the universe; there is no 'phlogiston', and spontaneous generation is impossible.

If those past mistaken conjectures prove one thing, it is that we *can* get it wrong. Some would take the argument a step further, however, in what (the early 'realist') Putnam (1978) labelled 'the disastrous meta-induction' of scientific reasoning. It infers, from the fact that those successful past scientific theories appear to be wrong, that our present theory must share the same fate and fail to refer accurately to the external world. Indeed, as Goldman (1986:158) writes: 'Suppose scientists decide that no term in the science of more than fifty years ago referred, and no theory of that vintage was true. If this keeps happening, shouldn't we be led to the meta-induction that no theoretical term now or in the future will refer, and no present or future theory will be true'?

This line of reasoning, however, I argue along with Goldman (1986:158-160), is severely misguided. First of all, Goldman claims, it is self-undermining. We can only judge past theories to be false in the light of present theories. If we derive from this that our present theories must be false, we lose all grounds to claim that those past theories are false. All we can claim, therefore, is that, while our past theories are definitely false in the light of present evidence, our present theories *could* be false in the light of future evidence.

Furthermore, defenders of the pessimistic meta-induction often focus on highly speculative theories in science. Quantum theory and relativity theory may, indeed, very well prove to be (partially) wrong at some point in the future, but it is unfair to restrict considerations to those types of theory. Indeed, the more 'prosaic' theories that livers detoxify the blood, that liquids boil at a certain temperature, and that water is composed out of hydrogen and oxygen, are not at risk of being radically overthrown somewhere in the future of scientific enquiry. Finally, theories have numerous components. Even when the theory as a whole is rejected, not all of the components are necessarily overthrown. The general theory of the electron, for instance, has changed over time, but the view of the electron as a basic unit of electrical charge, irresolvable into smaller units, has been retained (Goldman 1986:160). Similarly, the theory of evolution has been amended and completed since Darwin, but its basic premise – i.e. that all living organisms evolve by means of natural selection – is still valid, as is Copernicus's heliocentric model of the universe. In this regard, I argue, rather than concluding from past mistakes that we can never get it right, we can only infer that we might get it wrong.

We have to admit, therefore, to the possibility that our representations can, in principle, be mistaken, in other words, that they cannot, or can only poorly, fit an external object. While it is possible that our grasps of the world correspond to states of the external world, there is no absolute guarantee there actually is a correspondence. In this regard, while we can argue that there is epistemic access to the world, it would be wrong to see this correspondence as a 'given'. Our representations of the world do not, in other words, automatically refer to the external world. Therefore, inferring the existence of a corresponding external object from our mere possession of a mental representation is fallacious, and even contradicts the basic tenet of metaphysical or ontological realism - that there is a mind-independent reality. Indeed, equating the world to our representation of it reduces it to our mental realm, and sucks us into a Berkeleyan vortex of idealism.

This brings us to the conclusion that, while our representations can be about real objects and properties, they can also fail to represent anything. Our minds, indeed, are natural organs with fallible reasoning power, and the history of its products shows that this can and does occur more often than not. However, as argued above, this does not provide us with a reason to claim that no representations, present or future, will ever fit their objects. Imperfect as it is, the human mind is able to represent real objects and properties.

4.2 Limits set by computational capacity

As McGinn (1994:I) points out, we need to distinguish between two different potential sources of cognitive limitation. The first is related to what McGinn calls 'the content of our mental representations'[38], which refers to that which I have called 'the cognitive grasps' we have at our disposal. This concerns the 'range of concepts' we can deploy in thought or – in my terms – the representations that are included in our conceptual space. As argued above (cf. 2.1 Saving epistemological realism: the distinction between grasp and content), the fact that these representations are particular and contingent does not entail they cannot represent external objects.

The second potential source of cognitive limitation, however, is what McGinn (1994) calls 'the specific character of the operational system'. This refers to the processing abilities we have at our disposal. As McGinn points out, 'a system might be confined by what it can do with them [the concepts at its disposal] – say, because of attentional or memory limitations'

[38] This is not to be confused with Nagel's (1986) content of a representation, which points to what a representation refers to.

(I). There is, in other words, a limitation in computational capacity that could limit our ability to represent the world. Indeed, the fact that the human mind has limited processing power hardly needs to be argued for. We can only memorise so much in a certain period of time, can perform unassisted calculus only to a certain level of complexity, and can only play out a limited number of chess moves in our mind in a certain lapse of time.

Nevertheless, we are not restricted to the 'bare' mental powers we are endowed with in our attempts to represent the world. Our memory is dramatically extended by our access to libraries and other databases. Furthermore, our computational powers are equally increased by calculators, computers and the like. In this regard, the natural capacity of our minds is exponentially boosted by artificial aids. This, however, does not imply that there can be no limitations to our representational abilities imposed by the available computational capacity, including the cognitive aids we can or even possibly could bring to the task.

In this regard, computational complexity theory considers some problems 'intractable'. This entails that, while these problems can be solved in principle, they would require an infinite or astronomical amount of resources – such as, for instance, time, space, or memory – in order to be solved. When faced with an intractable problem, we do not, therefore, lack the proper computational operations to solve the problem, but rather the sufficient resources or capacity to carry out the operations. A famous example of intractability is given by the 'travelling salesman problem'. Menger (1930) describes it in the following words:

> 'It is the task to find, for finitely many points whose pairwise distances are known, the shortest route connecting the points. Of course, this problem is solvable by finitely many trials. Rules which would push the number of trials below the number of permutations of the given points, are not known. The rule that one first should go from the starting point to the closest point, then to the point closest to this, etc., in general does not yield the shortest route'.

Mathematicians, computer scientists, physicists and even chemists have mulled over the problem for over a century and systematically found large numbers of variables intractable (Appelgate et al, 2007:1). To date, the largest number of points successfully connected reached 85 900 points, using the powerful computer program ‘Concorde TSP solver’ (Appelgate et al, 2007:53). While increasing computational power and better heuristics can still increase this number, some magnitudes are bound to remain intractable.

While this is a purely mathematical problem, it is possible that some levels of complexity in which the physical world could theoretically be grasped might also elude us due to the restricted availability of cognitive resources. Indeed, if we assume that reality is infinitely complex - that there is, in other words, no external limit to the level of complexity in which the world can be represented, we have to admit that our grasp could never get to ‘the bottom’ of its enquiry. If, on the other hand, reality is finitely complex or – in other terms – if there exists a point in enquiry where the world would harbour no finer structures than those already represented, it is still not guaranteed that we could muster up the proper resources to reach this point. This computational limitation could surface in two different ways. We could either be unable to reach the ‘perfect’ theory of the world because at a certain level of complexity it would become intractable or – on the other hand – it could very well be that, even if we could reach that perfect ‘theory’, we would still be limited by our available conceptual capacity and so could not apply it to all present and future instances of the world to predict all present and future states of the world.

Some degrees of understanding and predicting, therefore, are very likely to remain un-reached, not because of their resistance to resolution by the computational tools at our disposal, but because of the amount of computational resources they would require in order to be reached. If, as argued above, there is no limit to the level of complexity in which the world can be grasped, no matter how much we increase the computational power at our disposal, our scope on the world will always be restricted to a certain level of resolution. If, on the other hand, there is such a limit, that limit could very well be beyond the computational resources we can bring to the task and even the resources we could theoretically (in the future) delve into.

4.3 Limits set by scope

As argued above (cf. 2.2 The nature of cognitive grasps: impossibility of a God's eye view), there is a one-way implication between content and grasp or – in Quine's terminology (1975:313) – between observation and hypothesis, meaning that the same content can in principle be yielded by different grasps (cf. 2.2 The nature of cognitive grasps: impossibility of a God's eye view). Every grasp, however, yields or 'implies' a singular content. Indeed, while a particular point in space can be grasped by both an axis-system and a grid-system, a point in the axis-system only refers to one point in space. Similarly, while a tree can be represented in a drawing or a description, both the drawing and the description only fit one instance in the world; namely, the tree they represent. Therefore, while an external object can be 'fitted' by a number of different grasps, a grasp only 'fits' a singular object.

This opens the possibility that all our available representations do not cover all the content the external world harbours. In other words, some properties or elements of the world might be inaccessible to us because no grasps or representations at our disposal can fit these properties or elements. This is not, however, as argued above, necessarily the case. A particular set of grasps can, in principle, yield a complete representation of the world. Indeed, not only could a particular grasp access the world in its totality, as I claimed, but such a complete representation would necessarily take a particular grasp (cf. 2.2 The nature of cognitive grasps: impossibility of a God's eye view).

Nevertheless, I suspect, some elements and properties of the world might elude us because of the nature of our grasp. They might, in other words, resist being grasped by the conceptual tools we have at our disposal. Indeed, either there is simply no representation in our conceptual space to fit certain external objects (closure) or we are (very) unlikely to form fitting representations, due to our particular cognitive nature (bias) (cf. Chapter 4). Recall, in this context, our numeracy lacking human beings (cf. Chapter 4 - 3.2.2 Cognitive resources). Not possessing the innate faculties underlying the ability to represent numerical information, they would – as pointed out – be hopelessly closed off to some fundamental ways of representing the world (in this case: through mathematics, and by extension, science as we know it). This raises serious doubts as to their ability to represent all the properties and elements of the external world we represent by means of their impaired grasp (impaired in comparison to us, that is).

In this regard, it is very unlikely that our numeracy lacking humans would be able to represent Einstein's space-time continuum (assuming this is an accurate representation of the world and therefore refers to an external property), or at least that they are biased against forming such a representation as they are lacking the concepts of mathematics and are therefore forced to take a different 'conceptual road' to arrive at a grasp which would represent this property of the world. Just as our a-numerical humans, lacking the innate faculty to represent numerical information, can be suspected to be most probably closed off to or at least biased against representing some aspects of the world, we could be lacking certain (theoretically possible) faculties, which would enable us to connect with a realm of external objects and properties that are now either straight-forwardly inaccessible to us (because no representation in our conceptual space fits these objects) or highly unlikely to be grasped by us (because we are biased against forming the necessary representations).

Furthermore, the possibility that some elements of the world fall outside our epistemic scope makes sense from an evolutionary perspective. As pointed out in chapters 2 and 3, our evolutionary past has endowed us with a set of representational tools (i.e. perceptual and cognitive abilities) that provide us with a picture of the world that enhanced our ancestor's chances for survival and successful reproduction. While we are not bound by this particular picture of the world, all our transgressing representations are still the product of recombining the perceptual and cognitive abilities that we are endowed with (cf. Chapter 3, 6.2 Variation through recombination). Considering that these tools only took shape when they provided a 'tangible' advantage in terms of biological fitness, it is very plausible that the combinatorial set of the representations generated by these biological tools – i.e. our conceptual space – do not fit or represent some aspects of the world. Anything else, while not impossible in principle, would be highly unlikely, in much the same way as the possibility that the parts and tools designed for building a car could be rearranged and redeployed to build an impeccable airplane. There is, indeed, no reason to suppose that tools designed for one purpose (i.e. a picture of the world boosting our biological fitness) are perfectly apt to fulfil another (i.e. representing the world theoretically or scientifically).

However, even if the scope of the representations within our conceptual space – i.e. all the representations we could possibly form – does not cover some elements and properties of the world – which therefore could never be grasped – this does not entail that representations we can possibly form cannot represent the world accurately, merely that they cannot represent it completely. Moreover, the incompleteness of our best possible theory or theories about the world is already strongly suggested by the fact that we are most probably limited, by the available computational capacity, to a certain level of complexity, inferior to the level of complexity in which the world could theoretically be grasped (cf. 4.2 Limits set by computational capacity).

Furthermore, I argue, this incompleteness – in the sense that some external objects resist being grasped by the representations we can possibly form, not in the sense that some levels of complexity are inaccessible due to limitations in computational capacity – would only be evident from a hypothetical external perspective, not from our own point of view. There would, in other words, be no 'black holes' in our best possible science: no irresolvable mysteries hovering in plain sight but all the while resisting satisfactory explanation. The ideal limit of human enquiry, in this regard, would provide us with a 'complete' and accurate representation of the world *for us*. The fact that some other cognitive creatures could assess that we fail to represent some aspects of the world does not make our grasp of the world either erroneous or inherently flawed. It merely entails that our cognitive relation to the world includes but a part of the elements and properties that could in theory be represented by a cognitive creature.

My claim that our ideal understanding would be void of any irresolvable mysteries is opposed to the 'New Mysterianism' movement, advocating 'Transcendental Naturalism'. The latter is championed by McGinn (1989, 1993, 1994). McGinn's argument contains two distinct claims. The first is that, due to the nature of our cognitive apparatus, we are closed to certain aspects of the world. The second is that we can pinpoint those aspects of the world that we are closed to. While I agree with the first claim – although not in principle – I will reject the second.

According to McGinn's (1989) Transcendental Naturalism, the human mind is cognitively closed to the solution of certain problems, 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': they transcend our cognitive capacities, but at the same time are not supernatural; 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. […] The mind-body problem is the problem of understanding how the miracle is wrought, thus removing the sense of deep mystery. We want to take the magic out of the link between consciousness and the brain (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 grasp 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, no different from problems we do solve) but because of the nature of our cognitive apparatus.

However, I object, along with Dennett (1991), that this line of thinking ignores one of the main characteristics of human thought: its productivity. As pointed out in chapter 4, we can produce an infinite amount of distinct and meaningful representations of the world (cf. Chapter 4, 2.1 Productivity of thought). This should, as Dennett claims, enable us to formulate and understand sentences or propositions that best express the solution to this so-called humanly irresolvable problem. There must, in other words, be representations within our conceptual space that best represent the psychophysical link.

Indeed, where McGinn's argument fails is not so much in arguing that we are cognitively closed to certain aspects of the world – which I argued is very likely, although it does not follow in principle from the particular grasp we have of the world – but in arguing that we are cognitively closed to the answers of problems we pose.[39] In this sense, when McGinn (1989:351) invokes the notion of monkeys being closed to the properties of electrons in order to argue that humans might very well be closed with respect to certain true theories, he overlooks the fact that the properties of electrons are not the answer to problems the monkey mind can conceive of. Indeed, unlike monkeys, we understand the mind-body problem. As Dennett (1991) points out, 'for McGinn to have a convincing case for human cognitive closure, he should provide an empirical example of some creature, human or otherwise, who can definitely understand some question, but be definitively incapable of understanding the answer'.

In this regard, I argue that every intelligible problem would have to have a potential ideal answer equally intelligible to us. If, however, that answer falls short of the objective reality it attempts to grasp, then that very question would fall equally short of describing the objective problem it points at. It does not seem possible that we could understand a problem in all its aspects and not be able to understand an equally 'complex' ideal answer. Our cognitive limitations would have an effect both on question and answer. It would be like someone who understands the concepts of numeracy and addition still not being able to solve the mathematical problem: x = 5+8. Therefore, if McGinn is right in claiming that some philosophical conundrums can never be satisfactorily answered, because we do not have the proper cognitive grasp at our disposal to solve the problem, how can we expect to have the

[39] To be fair, McGinn (1989:351-352) does address this issue, arguing that a hypothetical 'Humean mind', only able to derive concepts from perception, would be closed to a true theory of physics (unable to conceptualise the existence of unobservables as atoms) *and* baffled by the physical world because of this inability to form an accurate theory. This claim, however, is not supported by any argument, but merely offered as a fact.

proper grasp to understand the problem? The representations we use to describe the problem would have to be equally defective (our understanding of physical and mental categories and of interaction or nexus in the case of the mind-body problem) and therefore the very problem we pose would be, to say the least, a very poor representation of what really is at issue.

Delving into Boden's (1990) terminology of 'conceptual spaces' once again, we have – as argued in the previous chapter, an infinite amount of representations at our disposal – in other words, an infinite amount of locations in our conceptual space (cf. Chapter 4, 2.4 Conceptual spaces: reconciling openness and closure). While some aspects of the world might be ungraspable for us, because no representation in our conceptual space can possibly grasp them, those aspects can never be the object of our enquiry. Indeed, if on the one hand, we have an accurate grasp of the problem, there would have to be an ideal grasp of the answer in our conceptual space, since our conceptual space contains an infinite amount of representations and contains the proper representational tools to represent the issue at hand. If, on the other hand, an aspect of the world cannot be grasped by means of the representations within our conceptual space, we would be closed both to the question and to the answer. For instance, if a problem X can be grasped mathematically, there must be a correct mathematical solution to it. Therefore, to grasp the problem, one needs mathematics, the possession of which entails that a correct answer is within reach (by means of this same mathematical reasoning). If, however, some cogniser cannot engage in mathematical reasoning, it could never grasp the problem in the first place.

Therefore, while our scope on the world might not coincide with the world itself (which already is the case, assuming that we are limited by our available computational capacity); within this scope, I argue, there are no irresolvable 'black holes' in principle. This follows from our previously outlined characterisation of human cognition as an open cognitive system within a particular conceptual space. Whatever can be grasped by means of our conceptual building blocks can potentially be 'ideally' represented. Whatever cannot be grasped – i.e. that which falls outside our conceptual space – can never be an object of enquiry.

Furthermore, this view accords with the two-sided determination of knowledge I outlined above (cf. 3.1 The two-sided determination of knowledge), claiming that knowledge only makes sense within the particular cognitive relation we entertain with the world. In this regard, while elements of the world might fall outside this cognitive relation, there must be, within the scope of this relation, a theoretically ideal fit for every object of representation. Knowledge of the world – i.e. representations fitting the properties of the objects represented in virtue of our epistemic orientation – is, therefore, possible, and – moreover – the ideal limit of human knowledge, or 'the perfect fit', would be neither arbitrary (cf. biological determination), nor inherently flawed (there would be no irresolvable mysteries).

5. Conclusion

Gauging the implications of our 'open cognitive relation to the world, comprised within a particular conceptual space' for our epistemic endeavours, I argue that while we represent the world in a particular set of ways, this does not entail in principle that we can only represent it in a limited way. However, the way our representations correspond to the world is not in virtue of some objective 'mirroring', but in virtue of a particular way of fitting (i.e. our particular epistemic orientation) with a particular set of representational tools. Our knowledge of the world – i.e. representations corresponding to external objects – is, in this regard, fundamentally grounded in our biology. It is the outcome of our particular cognitive orientation, fuelled by our particular cognitive means. It is radically species-specific.

This, however, rather than sucking us into a bottomless vortex of relativism, bringing about scepticism, enables us – on the contrary – to resist pervasive and destructive forms of relativism. It provides us, indeed, with a necessary framework through which we *must* view the world (i.e. a framework rooted in our cognitive nature), rather than a multitude of contingent cultural frameworks. Knowledge, in this regard, while on the one hand losing its untenable universality – which I argue, moreover, is not the product of our limited cognitive means but follows necessarily from the very act of representing – gains, on the other hand, a 'human objectivity'. It is not a random, nor a conventional perspective on the world, but a necessary one in virtue of the kind of cognisers we are. It is the product both of the way the world is and the way we must represent it.

This brings my argument almost to a close. One more issue, however, needs to be addressed. That is the issue of the circularity of my approach – taking a particular human theory, i.e. the theory of evolution, to bear on the fundamental question as to what the scope and limits of human knowledge are – and its consequences. Indeed, it can be argued, how can we make claims about knowledge in general based on particular theories within our knowledge? This will be the subject of the next chapter.

Chapter 6: Evolutionary epistemology's double-edged sword

1. Introduction

In this dissertation, I have based epistemological claims on evolutionary considerations. The question now arises as to whether this is a coherent strategy at all. In this chapter I will argue that it is, defending my approach against what De Cruz et al (2011:518) call the 'double-edged sword' that evolutionary approaches to epistemology present us with. In this regard, I will claim that my argument is neither self-defeating, nor (viciously) circular.

2. Evolutionary epistemology's double threat

As De Cruz et al (2011) point out, evolutionary considerations can either be used to justify or to debunk beliefs. In other words, the premise that biological evolutionary forces have shaped our perceptual and cognitive faculties has been used to argue both that our beliefs about the world are justified, or – to the contrary – that those beliefs are unreliable. Indeed, as pointed out in chapter 2, the Lorenzian epistemologist argues that, since our perceptual and cognitive apparatus, which provide us with a representation of the world, are the result of evolution by natural selection, the structures of this representation must – at least approximately – match the structures of the world itself, because being endowed with an accurate representation of the world increases one's biological fitness, and therefore the perceptual and cognitive mechanisms allowing for this accurate representation must have been selected for (cf. Chapter 2, 2. Lorenz's evolutionary epistemology).

Opposed to this view is the claim that I defended with respect to our commonsense representations, which holds, contrary to the Lorenzian argument, that the fact that our perceptual and cognitive faculties evolved does not entail that the representations they cause us to have are (even approximately) accurate. Indeed, natural selection can only be expected to shape representations of the world that lead to fitness-enhancing behaviour. As Stich (1990) points out, these faculties will therefore not automatically yield accurate representations. Furthermore, natural selection is path-dependent, and evolution is also driven by other 'random' forces, namely genetic drift and the process of genetic hitch-hiking (cf. Chapter 2, 3. Uncovering the fallacy).

Both views – i.e. evolutionary justification arguments and evolutionary debunking arguments – run into problems. The former, as it relies on a theory which is the product of human belief-forming to justify the accuracy of human belief-forming, is inherently circular. Indeed, as Shogenji (2000) points out, 'it is commonly held by epistemologists that we cannot establish the reliability of a belief-forming process with the use of beliefs that are obtained by that very process since such self-dependent justification is circular' (287). Such an approach does, indeed, beg the question (i.e. commit the logical fallacy of 'petitio principii'), by assuming as a premise what is to be proven.

Evolutionary debunking arguments, which conclude that our representations are not reliable, are, on the other hand – as De Cruz et al (2011:517) point out – self-defeating. Arguing from an evolutionary standpoint that our cognitive faculties are unreliable, robs our premise (i.e. that our cognitive faculties evolved) of its reliability, since it too is a product of our 'non truth-tracking' cognitive faculties. This line of reasoning is at the core of Plantinga's (1993) 'evolutionary argument against naturalism', which claims that evolutionary naturalism is self-defeating, since, if we accept that our cognitive faculties are the product of evolution, we must accept they do not endow us with reliable beliefs (cf. Chapter 2, 3.1 Plantinga's case against evolutionary naturalism). Therefore, he argues, we have an 'undefeated defeater' of evolutionary naturalism, since the very idea that our cognitive faculties are the product of evolution then becomes unreliable. In other words, if we believe that evolution by natural selection has shaped our cognitive apparatus, we have good reasons to doubt the reliability of our beliefs, including the reliability of the belief that evolution by natural selection has shaped our cognitive apparatus.

In this regard, De Cruz et al (2011) argue that 'if evolutionary approaches to the human mind are to be coherent, they should allow at the very least for cognitive capacities that are capable of generating truth-tracking theories, such as evolutionary theory' (526). Indeed, if we eradicate the reliability of our cognitive processes, we completely undercut the grounding on which we base this conclusion. Granting cognitive reliability from evolutionary considerations, however, presents us – as pointed out – with the threat of circularity, supporting the reliability of our beliefs on the basis of a particular belief (which is therefore still in need of justification). In the next section, I will analyse how the argument presented in this dissertation fares in the light of this double threat.

3. Answering the double threat

The argument I have presented has three major parts. The first part derives from the premise that our perceptual and cognitive abilities are the product of a blind evolutionary process, that these abilities are contingent, and that they could have been different, therefore, providing us with a different perceptual input and a different cognitive processing of this input (cf. Chapter 1 – Chapter 2). The second part states that, while we are not bound by the commonsense representations these particular abilities provide us with (i.e. we are able to transgress these representations), we can nevertheless only produce representations that are included in the conceptual space defined by the combinatorial possibilities of the various representational tools we possess (cf. Chapter 3 – Chapter 4). Finally, the third part states that the way in which we fit representations to the external world is by means of our biologically determined epistemic orientation (cf. Chapter 5). Therefore, I concluded, we entertain a particular cognitive relation to the world, being endowed with a particular set of perceptual and cognitive resources, guided by a particular epistemic orientation. An accurate representation for us, in other words, is a particular fit (epistemic orientation) with particular means (perceptual and cognitive resources) (cf. Chapter 5).

3.1 Is my argument self-defeating?

Is this argument self-defeating? The self-defeating character of an argument can either be complete or not. When it is complete, the conclusion of the argument entails *necessarily* that the premise on which it is based is false. When it is not completely self-defeating, the conclusion of the argument entails that the premise, while not false in principle, is nevertheless unreliable – i.e. likely to be false. The first, completely self-defeating argument, I will call 'self-refuting'; the second, weaker form of a self-defeating argument, I will call 'self-undermining'.

Is our argument self-refuting? An evolutionary debunking argument would be self-refuting if it infers, from the premise that our perceptual and cognitive faculties evolved, that all the beliefs or representations they generate are necessarily false. The premise on which this conclusion is based, therefore – since it is itself a representation generated by our perceptual and cognitive faculties – must be false. The argument, in this case, refutes itself, since if its conclusion is true, its premise must be false.

The argument presented in this dissertation, however, is not self-refuting. Firstly, the fact that our perceptual and cognitive faculties are the outcome of a blind evolutionary process does not entail that all representations they cause us to have are necessarily false, but merely that these representations do not necessarily represent the world accurately *because* they originate from evolved faculties (cf. Chapter 2). Secondly, as argued in the fourth chapter, we can transgress our commonsense representations, substituting them with representations that we perceive as epistemically preferable. In this regard, even if evolution only provides us with false (commonsense) beliefs (which I think it does not), we would still have the cognitive means to discard these beliefs and replace them with other (better fitting) representations.

This leaves the weaker possibility that our argument is self-undermining. This would be the case if I inferred from my premises (cf. above) that our beliefs, while not necessarily false, are nevertheless unreliable – unreliable in the sense that the probability that any of our representations are accurate representations of the world, or, in terms of the last chapter, the probability that they fit the external world, is low. If I were to reach such a conclusion – i.e. that our representations of the world are unreliable – it would be very improbable that the premises on which this conclusion is based are accurate. Therefore, while the conclusion does not refute its premises (i.e. necessarily entails that they are false), it nevertheless undermines them.

The conclusion I reached, however, does not state that our representations of the world are unreliable. While this might be true to a certain extent for the commonsense representations we have of the world (cf. Chapter 1 and 3), our ability to transgress these representations entails that we are able to produce representations we perceive as better fits. Our ability to transgress, therefore, increases the reliability of our beliefs. It substitutes our (largely) unreliable commonsense representations with other representations, which are selected precisely because they are better fits and, therefore, more reliable.

In this regard, note that the theory of evolution on which my argument is based is a transgression of our commonsense representations of the organic world. Indeed, as pointed out in Chapter 1, at a commonsensical level we divide the organic world into immutable categories and sub-categories based on an intuition of a hidden trait or essence that members of the same group share with one another (cf. Chapter 1, 3.2.2 Folk biology). The reliability of my premises, in this regard, can be considered to be high (or at least, higher than those of our commonsense representations), since they are the product of a 'critical fit' to the input we gather from the world.

Another threat, however, could come from the fact that we entertain a particular cognitive relation with the world, fitting our representations to the particular and contingent epistemic orientation with which we are endowed (cf. Chapter 5), entailing that fitting representations are only fitting *for us*. This, however, does not undermine our premises. Our premises can indeed, as pointed out, fit the external world for us and therefore be accurate for us. This entails that they can be reliable to us. The conclusions we draw from premises which are reliable to us can, in this regard, also be expected to be reliable to us. The fact that other cognitive beings could produce different representations of the matter (the product of different cognitive tools and a different orientation), and that different representations would therefore be reliable to them, doesn't entail that our fitting representations are not reliable, or that they are less reliable, but merely that the reliability of a belief is determined by the epistemic relation a cognitive being entertains with the world (cf. Chapter 5).

A final threat comes from the fallibility of our representational abilities (cf. Chapter 5 – 4.1 Limits set by fallibility). In this regard, our current theory of evolution could be mistaken in some important respects (i.e. it could only poorly fit what is the case). However, even if this is the case, the fact that our perceptual and cognitive abilities are particular and contingent would still stand. The accuracy of our premise, therefore, is not highly dependent on the degree of accuracy of our best theory of how our cognitive and perceptual abilities were shaped, but merely on the less problematic assumption that these abilities are contingent products of a natural process.

3.2 Is my argument circular?

Evolutionary based epistemologies, however, present us with a 'double-edged sword', as De Cruz et al (2011:518) so eloquently phrase it. While we might have averted the threat of the self-defeating character of the argument presented here, we still face the threat of circularity. Is our argument circular? In order to answer this, we need to distinguish between two different threats – on the one hand, the threat of circularity with respect to the particular argument presented, and, on the other hand, the threat of the circularity of our approach in general (i.e. the alleged circularity of a naturalised form of epistemology). The former will be discussed here, while the latter deserves a new section altogether (cf. 4. Defence of a naturalised epistemology).

With regards to the particular argument presented in this dissertation, I argue that it is not circular. Indeed, I do not – as opposed to the Lorenzian epistemologist (cf. Chapter 2, 2. Lorenz's evolutionary epistemology) – ground our ability to represent the world accurately in the fact that our cognitive abilities evolved, which commits the fallacy of grounding the reliability of our beliefs in a particular belief. In the contrary, I argue that the fact that these faculties evolved does not guarantee the accuracy of the representations they provide us with (cf. Chapter 2). It is precisely the ability to transgress those representations in the light of our cognitive orientation that enables us to tune our representations to the world; in other words, to produce a cognitive fit.

Simply put, the Lorenzian epistemologist bases the accuracy of our representations on evolution. I, on the other hand, base the reliability of our ability to represent the world in our possession of an open epistemic relation to the world, guided by an epistemic orientation – i.e.

in our ability to transgress (cf. Chapter 3, Chapter 4 and Chapter 5). Is this circular? No. Whereas the Lorenzian argument bases the reliability of our representations on one representation (i.e. the theory of evolution), I do not base the reliability of our representations on a particular representation we have of the world. I base it on the epistemic relation we entertain with the world.

Furthermore, rather than concluding the (approximate) accuracy of our beliefs, as the Lorenzian epistemologist does, I conclude that we fit representations to the external world by means of a particular set of tools in the light of a particular epistemic orientation. If anything, therefore, my conclusion casts the accuracy of my premises in doubt, since it holds that accurate representations for us are particular fits by particular means, bringing about the self-undermining problem I addressed above. This threat, however, can, as argued above, be satisfactorily averted.

This leaves the threat of the circularity of my approach in general. Any form of naturalised or descriptive (i.e. non-normative) epistemology, can, indeed, be said to be circular, since it brings empirical findings (i.e. instances of knowledge) to bear on a theory of knowledge (i.e. knowledge in general), thereby presupposing what must be justified. This will be the subject of the next section.

4. Defence of a naturalised epistemology

Naturalised forms of epistemology make use of empirical data (i.e. scientific theories) to reflect on epistemological questions (i.e. questions concerned with the status, scope and limitations of knowledge). These forms of epistemology, therefore, make use of instances of knowledge (scientific theories) to reason about knowledge in general. Many philosophers have criticised such an approach to epistemology, arguing that it is hopelessly circular. Epistemology, according to these thinkers, must be a purely a priori analysis of the conditions which make knowledge possible, providing the framework to which ‘positive’ claims to knowledge have to yield. In the following pages, I will answer this criticism on the basis of both the argument that no form of epistemology can possible escape circularity, and that all forms of circularity or self-dependence need not be vicious, closed circles.

4.1 No alternative

The problem of circularity is not restricted to evolutionary epistemology or other kinds of naturalised (i.e. empirically informed) epistemology alone. It is at the core of any theory of knowledge. A theory of knowledge does, indeed, apply knowledge to reflect on knowledge. In providing our epistemic endeavours with foundations, we do so on the basis of criteria which are themselves in need of foundations. In this regard, Vollmer (1987) points out that an investigation into the validity of knowledge presupposes a criterion determining whether and when knowledge is valid – demarcating, in other words, genuine or true knowledge from beliefs or representations lacking such justification (164).

This criterion, Vollmer (1987) continues, can either be a piece of knowledge itself, or not. If it is a piece of knowledge, we are caught in a circle or an infinite regress. Indeed, either our criterion of validity already presupposes what it must determine (i.e. knowledge) or it is itself in need of a foundation, which in turn will require a foundation, ad infinitum (164). If, on the other hand, it is not a piece of knowledge, it would be an axiom or a convention. This, however, as Vollmer points out, could never justify knowledge, since such a justification would have to be recognisable as a criterion of objective knowledge. A criterion of knowledge, therefore, has to be a piece of knowledge itself – i.e. we need epistemic reasons to consider it as a criterion of objective knowledge. This, however, brings us back either to a circle or to an infinite regress (164).

This argument, Vollmer (1987) clarifies, does not entail that knowledge is impossible, but merely that an absolute justificatory theory of knowledge is impossible, in the sense that it cannot escape Munchausen's trilemna, meaning that it is either caught in a circle or an infinite regress, or that it is validated by an axiom. It is the problem of knowledge of knowledge (165). In this regard, 'If epistemology is indeed meant to justify knowledge, to produce sufficient criteria of truth, to establish necessary propositions, then it is useless, barren, impossible, stillborn. If we set out doubting the validity of every piece of knowledge, no knowledge can be justified' (Vollmer, 1987:165). In other words, a theory of knowledge cannot start from nothing. As Russell argues:

> If we adopt the attitude of the complete sceptic, placing ourselves wholly outside all knowledge, and asking, from this outside position, to be compelled to return within the circle of knowledge, we are demanding what is impossible, and our scepticism can never be refuted. For all refutation must begin with some piece of knowledge which the disputants share; from blank doubt, no argument can begin. Hence the criticism of knowledge which philosophy employs must not be of this destructive kind, if any result is to be achieved. Against this absolute scepticism, no logical argument can be advanced (Russell, 1912:112).

Epistemology, in this regard, cannot 'take off' without presupposing knowledge. We must, as Neurath's (1959) famous metaphor suggests, rebuild our ship while staying afloat, unable to start from 'scratch', or, as Neurath puts it, 'never able to dismantle it in dry-dock and to reconstruct it there out of the best materials' (201). Knowledge of knowledge, however, cannot escape a certain circularity.

Nevertheless, as Vollmer (1987:165) argues, the fact that epistemology cannot meet the 'excessive and self-contradictory requirements' that were traditionally attributed to it (i.e. to provide a criterion which is both non-arbitrary and not in need of further foundation) does not mean it should be abandoned. We have to change our conception of epistemology. It does not 'prove the existence of knowledge, it presupposes knowledge'(165). In this regard, Vollmer concludes, 'It [epistemology] is neither infallible nor unfailing and should not claim to be. It works hypothetico-deductively, as any other scientific discipline does' (165).

Epistemology, therefore, has no alternative but to lose its aspirations to be a 'first philosophy', attempting in Cartesian fashion to construct knowledge from the bottom up, and presupposing no knowledge when laying the foundations on which all knowledge is to be erected. As Quine (1969), argues, this traditional epistemological project has turned out to be a failure, and we are consequently left with no other alternative than a naturalised form of epistemology, reflecting on epistemological issues from a background of established beliefs. There is, therefore, no escaping a certain form of circularity or self-dependence in our attempts to formulate a theory of knowledge. This self-dependence, however, need not be vicious.

4.2 Virtuous circularity

The fact that we are deprived of absolute epistemological foundations does not mean we have to give up epistemology altogether. We can still reflect on human knowledge from premises which are not absolutely certain. This, as pointed out above, makes the approach circular or self-dependent, since the conjectures on which we base our epistemological considerations are themselves 'pieces of knowledge' (and therefore in need of justification). Circularity, nevertheless, is not vicious per se. As Buskes (1998) points out: 'As long as our presuppositions and premises are not sacrosanct but instead open to revision and refinement we are not caught in a vicious circle at all' (13). In this regard, by granting our premises a hypothetical status, enabling us to critically examine them and revise them, if necessary, over time, we escape a closed circle and enter what Vollmer (1987:179) calls 'an open spiral or self-correcting feedback loop'. In this context, Vollmer argues, that while 'of course, we cannot reasonably call in question our premises all at once, we should be ready to examine each of them in due time' (179). Vollmer explains:

> It [the process of a self-correcting feedback loop] is a continuous interplay, a perpetual give and take, a never-ending critical co-operation, an endless mutual correction. This process is not circular. We might rather liken it to a spiral. This spiral structure obtains both historically and systematically: New knowledge has necessitated new epistemologies, and new epistemological concepts have helped the advancement and the understanding of scientific theories (Vollmer, 1987:183).

Therefore, Vollmer concludes:

> The circle supposed to occur in empirically-oriented epistemologies is a virtuous circle [i.e. a circle which is not only consistent, but also fruitful and productive]. If there is any circularity at all, it is not a *petitio principii*, because there is no principle to be begged. Evolutionary epistemology is a hypothetico-deductive system which starts from a combination of factual and epistemological premises, trying to draw

> conclusions, to check them for consistency and for truth, and to correct them if necessary (Vollmer, 1987:182, his italics).

In a similar vein, Shimony (1981:101) claims that scientific results can shed light on the reliability of human cognition and – reciprocally – that these considerations can give adequate justification to the process of scientific investigation. This 'dialectical framework', as Shimony calls it, accepts tentative suppositions at the beginning of enquiry, which are then subject to criticism, and may be revised and refined. 'This dialectic', Shimony therefore concludes, 'is open, with no foregone conclusions and no suppositions that are so entrenched that they cannot be critically evaluated' (101).

The alleged circularity of naturalised epistemology, in this regard, can be viewed as an epistemic upward spiral, in which scientific conjectures lead to epistemological considerations, which in turn refine the scientific conjectures, leading to better epistemological criteria and methods, and so on. It enables us to improve our epistemic situation, all the while acknowledging the necessity to reason from somewhere – i.e. to take particular 'pieces of knowledge' (void of absolute certainty) as the starting point of epistemological considerations.

Indeed, knowledge, as Popper (1957) rightly argues, proceeds by conjectures and refutations, and it is precisely our willingness to critically reflect on the conjectures we take as premises of our epistemological considerations, that demarcates what Popper calls 'critical thinking' from 'dogmatic thinking'. This critical thinking, or refusal to grant our premises a 'sacrosanct' status, as Buskes (1998:13) puts it, is exactly what shields empirically informed epistemologies from a vicious kind of circularity, in which premise and conclusion are trapped in a closed circle. While self-dependent to a certain extent, therefore, the interplay between (empirical) premises and (epistemological) conclusions becomes dynamic, as each feedback loop provides us with better conjectures.

In this regard the (epistemological) conclusions of the argument presented in this dissertation are not founded on a rusted set of dogmatic premises, but rather on a set of tentatively formulated hypotheses – i.e. our best conjectures to date. As argued above, these premises could very well be amended to an important extent in the future of scientific enquiry. This, in turn, would sharpen the emerging epistemological picture. The fact that these premises are not the absolute, immutable foundations the Cartesian epistemologists seek for their epistemic building, does not entail that the considerations on the status, scope and limits of human knowledge presented here, can and should be discarded as mere 'unfounded' ramblings. Rather, they should be viewed as a temporary picture to be enhanced in the future of enquiry.

In this regard, (future) research in evolutionary biology in general and cognitive and neuroscience in particular could provide valuable insights into how exactly human thought processes work; how, in other words, we form thoughts based on the resources that evolution has provided us with. This, in turn, would further define the epistemological issues of the scope and limits of human knowledge. It is a continuous process, bringing forth ever improving cognitive fits. In this context, as Einstein pointed out, 'there could be no fairer destiny for any theory than that it should point the way to a more comprehensive theory in which it lives on' (cf. Popper, 1957).

5. Conclusion

While an evolutionary informed epistemology faces the double threat of being either self-defeating or circular, the argument presented in this dissertation resists this 'double-edged sword'. Indeed, as argued, it is not self-defeating, since the evolutionary origin of our perceptual and cognitive apparatus does not impose on us the conclusion that our representations of the world cannot be trusted to be reliable. We are, indeed, able to transgress the commonsense representations – which are shaped in the light of biological fitness and are therefore not truth-tracking per se – by other representations that we perceive as better fits to the external world in virtue of our epistemic orientation.

Neither is our argument circular, since – on the one hand – I do not derive from the evolutionary origin of our cognitive apparatus that we are endowed with an accurate representation of the world, thereby grounding the validity of our beliefs in a single belief. On the other hand, I argued that our approach – basing epistemological considerations on empirical theories, i.e. basing a theory of knowledge on instances of knowledge – is, while circular, not viciously so. This self-dependence is, indeed, a critical one, where premises are not dogmatically accepted but subject to critical evaluation. The picture emerging from these premises is therefore, while not absolute, nevertheless justified, and constantly adapted to, or rather sharpened by, the best conjectures we possess.

Conclusion

The aim of this thesis was to fundamentally rethink our epistemic situation (i.e. the status, scope and limits of human knowledge), given the fact that the perceptual and cognitive abilities we are endowed with are shaped by the blind process of evolution by natural selection. In order to do so, I analysed our uncritical scope on the world determined by our perceptual and cognitive abilities (our commonsense theories) and the way we, nevertheless, overcome this 'myopic' and contingent, species-specific scope, by transgressing our biologically determined view of the world.

This transgression, however, is not boundless. Indeed, being the product of our perceptual and cognitive faculties, it is restricted to the conceptual space defined by the combinatorial possibilities generated by these faculties. This brings to light, a number of limits to our ability to represent the world. Furthermore, as transgression is guided by an epistemic orientation, human knowledge itself – i.e. the 'fit' between a representation and the world – is determined by this orientation, and therefore defined by the epistemic relation we hold to the world. It becomes species-specific – a product of our very own, biologically determined epistemic project.

Evolution, therefore, rather than 'simply' endowing us with accurate representations (as the Lorenzian epistemologists claim) or unreliable representations of the world (as the opposite sceptical camp claims), has endowed us with a particular set of species-specific perceptual and cognitive abilities, providing us with a set of contingent commonsense theories *and* the ability to overcome these representations, by digging into a conceptual space defined by the combinatorial possibilities generated by our perceptual and cognitive abilities. This provides us with an open but particular epistemic relation, enabling us to fit representations to the world in the light of our epistemic orientation. In this regard, while all other creatures we are acquainted with are restricted to a particular, 'given' scope on the world, homo sapiens radically broadens its cognitive horizons, being endowed both with the means to go beyond this scope (variation through recombination), and with the orientation to guide these explorations.

The cognitive abilities to transgress our biologically based representations, therefore, provide us with a new level of cognition altogether. Whereas other animal species possess a given cognitive framework in which they integrate the input they gather from their environment, we have the means to adapt this framework, i.e. we can override and replace the core intuitions we hold about the world and which make sense of the input data we gather. Language, as briefly pointed out (cf. Chapter 4) may have played a crucial role in this, since both our ability to metarepresent and our ability to map across domains and reason by analogy, seem closely tied to the development of natural language in the evolutionary history of humankind. Another crucial aspect of human cognition, is, however – and I cannot emphasise this enough – the possession of an epistemic orientation. Together, these distinctive features of the human mind underlie the remarkable achievements of human epistemic endeavours, resulting from this new level of cognition that humankind developed.

Despite its distinctiveness, however, human knowledge is still fuelled by a set of contingent perceptual and cognitive resources. This tension is at the core of the human epistemic situation. We transgress our biological view of the world, but can only do so based on our 'natural' abilities; we can produce an unlimited amount of different representations, but these representations are nevertheless comprised within a particular conceptual space. Human knowledge, in this regard, is neither an instance of God-like or universal rationality as so often assumed in the pre-Darwinian era, nor is it, at best, a boosted kind of primate cognition (as some evolutionary debunking arguments seem to suggest). The human mind is a natural product of evolution with certain limits, but it is nevertheless a unique feature in the evolutionary tree, endowing us, as pointed out, with a whole new level of cognition.

The knowledge resulting from this distinctive epistemic relation we entertain with the world, is an 'active' or conscious fit to the world by means of a contingent set of available perceptual and cognitive resources - a fit, in other words, resulting from a conscious enquiry into the external world, aimed at producing matching representations, and guided by an epistemic orientation. In this regard, while all other species possess a 'primary' cognitive relation to the world, endowed with a particular framework in which they integrate the input they gather from their environment, homo sapiens possesses a 'secondary' cognitive relation to the world, holding a framework by which he constructs frameworks which allow him to best represent the input he gathers from the world. Human knowledge, therefore, is not – as opposed to the representations of other animal species – a given perspective on the world fashioned by natural selection (natural selection, which is, in this case, the locus of the fit between representation and world) but a human creation, integrating the locus of this fit by virtue of an internal epistemic orientation.

This brings our discussion to a close. Reaching back to the question posed at the start of this dissertation: 'is knowledge of the world possible, given the fact that our perceptual and cognitive abilities are the product of evolution?', I answer emphatically: 'yes'. Knowledge, however, is not solely determined by the states of the world it represents, but also by the available perceptual and cognitive resources *and* the epistemic orientation we inherited from our evolutionary past. It is a maxim to which all human representations must adhere. It is neither simply given, nor forever unreachable, but a goal to be realised, firmly grounded in our cognitive architecture, shaped by the remarkably agile hand of the 'blind watchmaker' that is evolution.

As I am very well aware of, this dissertation – rather than 'closing' the debate – points the way to future research. For instance, I purposefully avoided committing to a theory of truth. The explicit formulation of a theory of truth, would have taken us far beyond the scope of this work. In this regard, my main claims were negative, arguing – on the one hand – that conceptualising correspondence by means of the mirror metaphor is untenable, while – on the other hand – at the extreme opposite, outright and radical anti-realistic standpoints are uncalled for.

Closely related to the issue of truth, is the debate about epistemological realism. Drawing a positive and elaborate view from the arguments presented here, comparing it with some of the important stances taken on the issue of epistemological realism - as Rorty's (1979) neopragmatism, Davidson's (1986) coherence theory of truth and Putnam's internal realism (1990)[40], to name but a few giants – we may well arrive at a novel position that could be fruitfully discussed in modern epistemology.

[40] A comparison beyond the tentative remarks made about Putnam and Rorty's views (cf. 2.2 The nature of cognitive grasps: impossibility of a God's eye view).

Finally, more than just pointing the way to further research in epistemology, the premises and approach of this thesis could find fertile ground in other areas of philosophy, such as morality. Indeed, the ability to override or transgress the innate set of implicit basic cognitive predispositions evolution provided us with, does not only hold for our representations of the world. It also holds for moral intuitions. In this regard, in the same way that we overcome our species-specific and uncritical view of the world, we overcome (in the sense of going beyond and against) our predisposed behavioural inclinations towards one another, other 'groups' and other species. In a similar way, a critical analysis of what underlies this feat, could point the way to the origin, scope and limits of moral theories. This dissertation, therefore, much rather than marking the end of an enquiry, marks its beginning. As so often in philosophy, every tentative answer points the way to more questions.

References

Appelgate, D., Bixby, R., Chvátal, V., Cook, W. 2007. *The traveling salesman problem.* Princeton: Princeton University Press

Artigas, M. 1992. *Human knowledge, reliability and fallibilism.* Pamplona: Universitad de Navarra, Group of Research on science, reason and faith (CRYF)

Atran, S. 1995. Causal constraints on categories and categorical constraints on biological reasoning across cultures, in Sperber, D., Premack, D., Premack, A. (eds.) *Causal cognition: A multidisciplinary debate.* Oxford: Clarendon Press

Atran, S. 1998. Folk biology and the anthropology of science: Cognitive universals and cul tural particulars. *Behavioral and brain sciences*, 21: 547-609

Ayer, A. 1959. *Logical Positivism*. Glencoe: The Free Press

Badcock, C. 2000. *Evolutionary psychology: A critical introduction.* Oxford: Polity Press

Baillargeon, R. 1991. Physical reasoning in infancy, in Gazzaniga, M. (ed.) *The cognitive neurosciences.* Cambridge: MIT Press

Baillargeon, R., Kotovsky, L., Needham, A. 1995. The acquisition of physical knowledge in infancy, in Sperber, D., Premack, D., Premack, A. (eds.) *Causal cognition: A multidisciplinary debate.* Oxford: Clarendon Press

Barbour, I. 1974. *Myths, models and paradigms: A comparative study in science and religion.* New York: Harper & Row Publishers

Barkow, J, Tooby, J, Cosmides, L. (eds.) 1992. *The adapted mind: evolutionary psychology and the generation of culture.* Oxford: Oxford University Press

Batteau, D. 1964. *The role of the pinna in human localization.* Cambridge: United Research Incorporated

Boas, F. 1887. Museums of Ethnology and their classification. *Science*, 9: 589

Boden, M. 1990. *The creative mind: Myths and mechanisms.* London: Weidenfeld and Nicolson

Boeckx, C. 2010. *Language in cognition: Uncovering mental structures and the rules behind them.* Malden: Wiley-Blackwell

Boyer, P. 2000. Natural epistemology or evolved metaphysics? Developmental evidence for early-developed, intuitive, category-specific, incomplete, and stubborn metaphysic al presumptions. *Philosophical psychology*, vol. 13, 3: 277-297

Boynton, R. 1979. *Human color vision.* New York: Holt, Rinehart and Winston

Bradie, M. 1986. Assessing evolutionary epistemology. *Biology & Philosophy*, 1: 401-459

Brewer, M. & Collins, B. (eds.) 1981. *Scientific inquiry and the social sciences: A volume in honor of Donald T. Campbell.* San Francisco: Jossey-Bass Publisher

Buskes, C. 1998. The Genealogy of Knowledge: A Darwinian approach to epistemology and philosophy of science. *Studies in general philosophy of science*, vol 4. Tilburg: Tilburg University Press

Buskes, C. 2006. *Evolutionair denken: De invloed van Darwin op ons wereldbeeld.* Amster dam: Uitgeverij Nieuwezijds

Campbell, B. 1985. *Human evolution: An introduction to man's adaptations.* 4th edition. New Yersey: Aldine Transaction

Campbell, D. 1974. Evolutionary epistemology, in Schlipp P. (ed.). *The philosophy of Karl. R. Popper.* LaSalle: Open Court Publishing Company: 412-463

Carey, S., Gelman, R. (eds.) 1991. *The epigenesis of mind.* Hillsday: Erlbaum

Carey, S., Spelke, E. 1994. Domain-specific knowledge and conceptual change, in Hirschfeld, L., Gelman, S. (eds.) *Mapping the mind: Domain specificity in cognition and culture.* Cambridge: Cambridge University Press

Carey, S. 2009. *The origin of concepts.* Oxford: Oxford University Press

Carroll, J. (ed.) 1956. *Language, Thought and Reality.* Cambridge: MIT press

Carruthers, P., Chamberlain, A. (eds.) 2000. *Evolution and the human mind: modularity, language, and metacognition.* Cambridge: Cambridge University Press

Carruthers, P. 2006. *The architecture of the mind: Massive modularity and the flexibility of thought.* Oxford: Clarendon Press

Chalmers, A. 1976. *What is this thing called science.* Cambridge: Hackett Publishing Company

Chomsky, N. 1959. Review of Skinner's Verbal Behavior. *Language*, 35: 26-58

Chomsky, N. 1965. *Aspects of the Theory of Syntax.* Cambridge: The MIT Press

Churchland, P. 1987. Epistemology in the age of neuroscience. *Journal of Philosophy*, vol. 84, 10: 544-553

Clark, A. 1986. Evolutionary epistemology and the scientific method. *Philosophica* 37: 151-162

Cosmides, L., Tooby, J. 2003. Evolutionary psychology: Theoretical foundations, in Nadel, L. (ed.) *Encyclopedia of Cognitive Science.* London: Macmillan

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

Darwin, C. 1871. *The descent of man and selection in relation to sex*. Detroit: Gale Research (1974).

Darwin, C. 1881. *Letter 3230 – Charles Darwin to William Graham.* July 3rd, 1881, retrieved from: http://darwinproject.ac.uk/entry-13230

Davidson, D. 1986, A Coherence Theory of Truth and Knowledge, in LePore, E. (ed.). *Truth And Interpretation, Perspectives on the Philosophy of Donald Davidson.* Oxford: Basil Blackwell.

Dawkins, R. 1986. *The blind watchmaker.* New York: Norton

De Cruz, H., De Smedt, J. 2007. The role of intuitive ontologies in scientific understanding – the case of human evolution. *Biology and Philosophy*, 22: 351-368

De Cruz, H., Boudry, M., De Smedt, J., Blancke, S. 2011. Evolutionary approaches to epistemic justification. *Dialectica*, vol. 65, 4: 517-535

de Vries, W. 2011. *Wilfrid Sellars.* The Stanford Encyclopedia of Philosophy (Summer 2011 Edition), Edward N. Zalta (ed.)

Dennett, D. 1991. Review of McGinn, The problem of consciousness. *The Times Literary Supplement*, May 10, 1991

Dennett, D. 1995. *Darwin's dangerous idea: Evolution and the meanings of life.* London: Allen Lane

Dennett, D. 2000. Making tools for thinking, in Sperber, D. (ed.). *Metarepresentations: A Multidisciplinary Perspective.* Oxford: Oxford University Press

Dercole, F., Rinaldi, S. 2008. *Analysis of evolutionary processes. The adaptive dynamics approach and its applications.* Princeton: Princeton University Press

Descartes, R. 1641. Méditations metaphysiques, in Adam, C, Tannery, P. (eds.). 1983. *Oeuvres De Descartes*, 11 vols., Paris: Librairie Philosophique J. Vrin

Dewey, J. 1910. *The influence of Darwin on philosophy and other essays.* New York: Henry Holt and Company

Diamond, J. 1992. *The rise and fall of the third chimpanzee. How our animal heritage affects the way we live.* London: Vintage

Dienes, Z., Perner, J. 1999) A theory of implicit and explicit knowledge. *Behavioral and brain sciences*, 22: 735-808

Dycke, R., Peebles, P., Roll, P., Wilkinson, D. 1965. Cosmic Black-Body Radiation. *Astrophysical Journal Letters*, 142: 414–419

Fodor, J. 1983. *The modularity of mind: An essay on faculty psychology*. Cambridge: MIT Press

Fodor, J. 1985. Précis of 'The modularity of mind. *The behavioral and brain sciences*, 8: 1-42

Galileo, G. 1623. Il Saggiatore, The Assayer. Translated by Stillman Drake in *The Controversy of the Comets of 1618*. Philadelphia: The University of Pennsylvania Press (1960)

Gelman, S., Wellman, H. 1991. Insides and essences. *Cognition*, 38: 214-244

Gelman, R., Durgin, F., Kaufman, L. 1995. Distinguishing between animates and inanimates: not by motion alone, in Sperber, D., Premack, D., Premack, A. (eds.). *Causal cognition: A multidisciplinary debate.* Oxford: Clarendon Press

Gentner, D. 2003. Why we're so smart, in: Gentner, D., Goldin-Meadow, S. (eds.). *Language in mind: Advances in the study of language and thought.* Cambridge: MIT Press

Gentner, D., Goldin-Meadow, S. (eds.). 2003. *Language in mind: Advances in the study of language and thought.* Cambridge: MIT Press

Gibson, J. 1979. *The ecological approach to visual perception.* Boston: Houghton Mifflin Company

Goldman, A. 1986. *Epistemology and cognition.* Cambridge: Harvard University Press

Granit, R. 1977. *The purposive brain.* Cambridge: The MIT Press

Haselton, M., Buss, D. 2000. Error management theory: a new perspective on biases in cross-sex mind reading. *Journal of personality and social psychology*, 78: 81-91

Haselton, M., Nettle, D. 2006. The paranoid optimist: an integrative evolutionary model of cognitive biases. *Journal of personality and social psychology review*, vol. 10, 1: 47-66

Hirschfeld, L., Gelman, S. (eds.) 1994. *Mapping the mind: Domain specificity in cognition and culture.* Cambridge: Cambridge University Press

Hoffmeyer, J. 1993. *Signs of meaning in the universe.* Translated by Haveland, B. 1996. Bloomington: Indiana University Press.

Hume, D. 1739. *A Treatise of Human Nature*. Oxford: Oxford University Press (1968)

James, W. 1890. *The principles of psychology.* Cambridge: Harvard University Press

Johnson, M. 2007. *The meaning of the body: Aesthetics of human understanding.* Chicago: University of Chicago Press.

Kahneman, D., Tversky, A. 1973. On the psychology of prediction. *Psychological Review*, 80: 237-251

Kahneman, D., Slovic, P., Tversky, A. (eds) 1982. *Judgment under uncertainty: heuristics and biases.* Cambridge: Cambridge University Press

Kant, I. 1781. Critique of pure reason, in Guyer, P., Wood, A. (eds.). 1998. *The Cambridge edition of the works of Immanuel Kant.* Cambridge: Cambridge University Press

Kaufman, D. 2002. Reality in common sense: Reflections on realism and anti-realism from a 'common sense naturalist' perspective. *Philosophical investigations*, 25: 331-361

Karmiloff-Smith, A. 1992. *Beyond modularity: A developmental perspective on cognitive science.* Cambridge: MIT Press

Kellman, P., Spelke, E. 1983. Perception of partly occluded objects in infancy. *Cognitive psychology*, 15: 483-524

Kemmerer, D. 2005. The spatial and temporal meanings of English prepositions can be independently impaired. *Neuropsychologia*, 43(5): 797–806

Khalidi, M. 2002. Nature and nurture in cognition. *British Journal of Philosophy of Science*, 53: 251-272

Kinzey, W. (Ed.). 1987. *The evolution of human behavior, primate models.* Albany: Suny Press

Kortlandt, A. 1965. How do chimpanzees use weapons when fighting leopards? *Year Book of The American Philosophical Society*, 5: 327-332

Kosslyn, S. 1980. *Image and mind.* Cambridge: Harvard University Press

Krellenstein, M. 1995. Unsolvable problems, visual imagery and explanatory satisfaction. *Journal of mind and behavior*, 16: 235-253

Kuhn, T. 1977. *The Essential Tension. Selected Studies in Scientific Tradition and Change*. Chicago: University of Chicago Press

Kukla, A. 2005. *Ineffability and philosophy.* New York: Routledge

Lakoff, G., Johnson, M. 1980. *Metaphors we live by.* Chicago: University of Chicago Press

Lakoff, G. 1987. *Women, fire, and dangerous things: What categories reveal about the mind.* Chicago: University of Chicago Press

Laudan, L. 1981. A confutation of convergent realism. *Philosophy of science*, 48: 19-49

Leeming, D., Leeming, M. 2009. *A dictionary of creation myths.* Oxford: Oxford University Press

Leibniz, G. 1704. New Essays on Human Understanding, in: Parkinson, G. *Leibniz: Philosophical Writings*. Translated by Morris, M & Parkinson, G. London:J.M. Dent & Sons

Levine, M., Shefner, J. 1991. *Fundamentals of sensation and perception.* Second Edition. California: Belmont, Brooks/Cole Publishing Company

Locke, J. 1690. Essays concerning human understanding, in Woolhouse, R. (ed.). 1997. London: Peguin Books

Lombrozo, T. 2007. Simplicity and probability in causal explanation. *Cognitive psychology*, 55: 232-257

Longino, H. 1990. *Science as social knowledge: values and objectivity in scientific enquiry.* Princeton: Princeton University Press

Lorenz, K. 1935. The companion in the bird's world: Fellow members of the species as releasers of social behavior. *Journal für Ornithology*, Beiblat 83: 137-213

Lorenz, K. 1941. Kant's doctrine of the a priori in the light of contemporary biology, in Plotkin, H. (ed.). 1982. *Learning, development, and culture. Essays in evolutionary epistemology.* New York: John Wiley & Sons

Lorenz, K. 1973. *Behind the mirror: a search for a natural history of human knowledge.* New York: Harcourt Brace Jovanich (1977)

Lumsden, C., Wilson, E. 1981. *Genes, mind and culture.* Cambridge: Harvard University Press

Malthus, T. 1798. *An essay on the principle of population.* London: Murray

Malotki, E. 1983. Hopi Time: A Linguistic Analysis of the Temporal Concepts in the Hopi Language, in: Winter, W. (ed.). *Trends in Linguistics. Studies and Monographs.* New York: Mouton Publishers

Marr, D. 1982. *Vision: A computational investigation into the human representation and processing of visual information.* New York: W.H. Freeman and Company

Matthen, M. 2005. *Seeing, doing, & knowing. A philosophical theory of sense perception.* Oxford: Oxford University Press

Mayr, E. 1950. Taxonomic categories in fossil hominids. *Cold Spring Harb. Symp. Quant. Biol.*, 15: 109-117

Mayr, E. 2001. *What evolution is.* London: Orion Books Ltd

McGinn, C. 1989. Can we solve the Mind-Body problem? *Mind*, New Series, vol. 98, 391: 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

Menger, K. 1930. Notes from a mathematics colloquium held in Vienna, February 5, 1930 [389]

Miller, A. 2000. *Insights of genius: Imagery and creativity in science and art.* Cambridge: MIT press

Miller, G. 2000. *The mating mind: how sexual choice shaped the evolution of human nature.* New York: Doubleday

Mithen, S. 1996. *The prehistory of the mind: The cognitive origins of art and science.* London: Thames and Hudson Ltd

Munz, P. 1993. *Philosophical Darwinism: On the origin of knowledge by means of natural selection.* London: Routledge

Nadel, L. (ed.). 2003. *Encyclopedia of Cognitive Science.* London: Macmillan

Nagel, T. 1986. *The view from nowhere*. Oxford: Oxford University Press

Neuhoff, J. 1998. Perceptual bias for rising tones. *Nature*, 395: 123–124

Neurath, O. 1959. Protocol Sentences, in: Ayer, A. *Logical Positivism*. Glencoe: The Free Press

Oeser, E. 1987. *Psychozoikum: Evolution und Mechanismus der Menschlichen Erkenntnisf ähigkeit.* Berlin: Parey

Papineau, D. 2000. The evolution of knowledge, in Carruthers, P., Chamberlain, A. (eds.). *Evolution and the human mind: modularity, language, and metacognition.* Cambridge: Cambridge University Press

Piattelli-Palmarini, M. 1994. *Inevitable illusions: how mistakes of reason rule our minds.* New York: John Wiley & Sons, Inc

Pinker, S. 1994. *The language instinct.* New York: Morrow.

Pinker, S. 1997. *How the mind works.* New York: W.W. Norton and company

Pinker, S. 2002. *The blank slate.* London: Allen Lane

Pinker, S. 2007. *The stuff of thought: Language as a window into human nature.* New York: Viking Pinguin

Plantinga, A. 1993. *Warrant and proper function.* New York: Oxford University Press

Plantinga, A. 1994. *Naturalism defeated.* Unpublished manuscript. Retrieved from: http://www.calvin.edu/academic/philosophy/virtual_library/articles/plantinga_alvin/naturalism_defeated.pdf

Plato 380 B.C.. Meno, in Jowett, B. (ed.). 1892. *The Dialogues of Plato translated into English with Analyses and Introductions by B. Jowett, M.A. in Five Volumes.* 3rd edition revised and corrected. Oxford: Oxford University Press

Plotkin, H. (ed.). 1982. *Learning, development, and culture. Essays in evolutionary epistemology.* New York: John Wiley & Sons

Popper, K. 1957. Philosophy of science: a personal report, in Mace, C. (ed.). *British Philosophy in Mid-Century.*

Popper, K. 1972. *Objective knowledge: An evolutionary approach.* Oxford: The Clarendon Press

Proffitt, D., Gilden, D. (1989). "Understanding natural dynamics." *Journal of experimental psychology: Learning, memory and cognition* 9: 384-393

Putnam, H. (1978). *Meaning and the moral sciences.* London, Routledge & Kegan Paul Ltd

Putnam, H. (1990). *Realism with a human face.* Edited by James Conant. Cambridge, Harvard University Press

Pütz, M., Verspoor, M. (eds.). 2000. *Explorations in linguistic relativity*. Amsterdam: John Benjamins Publishing Company

Quine, W. V. O. 1951. Two Dogmas of Empiricism, reprinted in *From a Logical Point of View*, 2nd Edition: 20-46. Cambridge: Harvard University Press

Quine, W.V.O. 1960. *Word and object.* Cambridge: Harvard University Press

Quine, W.V.O. 1969. *Ontological relativity.* New York: Columbia University Press

Quine, W.V.O. 1975. On Empirically Equivalent Systems of the World. *Erkenntnis*, 9: 313–328.

Radnitzky, G., Bartley, W. 1987. *Evolutionary Epistemology, rationality, and the sociology of knowledge.* La Salle: Open Court Publishing Company

Rescher, N. 1990) *A useful inheritance: Evolutionary aspects of the theory of knowledge.* Savage: Rowman & Littlefield

Ridley, M. 1993. *Evolution.* Cambridge: Blackwell Science

Riedl, R. 1984a. *Biology of knowledge.* New York: Willey

Riedl, R. 1984b. Evolution and evolutionary knowledge. On the correspondence between cognitive order and nature, in: Wuketits, F. (ed.). *C*oncepts and approaches in evolutionary epistemology. *Theory and decision library*, vol. 36. Dordrecht: D. Reidel Publishing Company.

Roediger, H., Craik, F. eds. 1989. *Varieties of memory and consciousness: Essays in honour of Endel Tulving.* Hillsdale: Lawrence Erlbaum Associates, Publishers

Rorty, R. 1979. *Philosophy and the mirror of nature.* Princeton: Princeton University Press

Ruse, M. 1986. *Taking Darwin seriousl*y. Oxford: Basil Blackwell

Ruse, M 1995. *Evolutionary naturalism. Selected essays.* New York: Routledge

Russell, B. 1912. *The problems of philosophy.* Los Angeles: Indo European Publishing (revised version of 1945, original publisher: Oxford University Press)

Schacter, D. 1989. On the relation between memory and consciousness: dissociable interactions and conscious experience, in: Roediger, H., Craik, F. (eds.). *Varieties of memory and consciousness: Essays in honour of Endel Tulving.* Hillsdale: Lawrence Erlbaum Associates, Publishers

Schlipp P. (ed.). 1974. *The philosophy of Karl. R. Popper.* LaSalle: Open Court Publishing Company

Sellars, W. 1956. Empiricism and the philosophy of mind, in Feigl, H., Scriven, M. (eds.). *Minnesota Studies in the Philosophy of Science*, vol. 1. Minneapolis: University of Minnesota Press

Shimony, A. 1981. Integral epistemology, in Brewer, M. & Collins, B. (eds.). *Scientific inquiry and the social sciences: A volume in honor of Donald T. Campbell. S*an Francisco: Jossey-Bass Publisher

Shogenji, T. 2000. Self-dependent justification without circularity. *Faculty publications*. Paper 218

Shopenhauer, A. 1819. *The World as Will and Representation*. Translated by Payne, E. 1969. New York: Dover Publications

Simpson, G. 1961. *Principles of animal taxonomy.* New York: Columbia University Press

Skinner, B. 1957. *Verbal behavior.* New York: Prentice Hall

Sober, E. 1981. The evolution of rationality. *Synthese*, 46: 95-120

Spelke, E. 1991. Physical knowledge in infancy: Reflections on Piaget's theory, in Carey, S., Gelman, R. (eds.). *The epigenesis of mind.* Hillsday: Erlbaum

Spelke, E. 2003. What makes us smart? Core knowledge and natural language, in Gentner, D., Goldin-Meadow, S. (eds.). *Language in mind: Advances in the study of language and thought.* Cambridge: MIT Press

Spelke, E., Phillips, A., Woodward, A. 1995. Infants' knowledge of object motion and human action, in: Sperber, D., Premack, D., Premack, A. (eds.).*Causal cognition: A multidisciplinary debate.* Oxford: Clarendon Press

Spencer , H. 1855. *The principles of psychology.* London: Longman, Brown, Green, and Longmans

Sperber, D., Premack, D., Premack, A. (eds.) 1995. *Causal cognition: A multidisciplinary debate.* Oxford: Clarendon Press

Sperber, D. (ed.) 2000. Metarepresentations: A Multidisciplinary Perspective. Oxford: Oxford University Press

Sperber, D. 2000. Metarepresentations in an evolutionary perspective, in: Sperber, D. (ed.). *Metarepresentations: A Multidisciplinary Perspective.* Oxford: Oxford University Press

St. Clair, R. 2007. The Sedimentation Theory of Cultural Time and Space: The Present is Embedded in the Past. *CÍRCULO de Lingüística Aplicada a la Comunicación*, 31: 52-90

Stanovich, K. 2004. Metarepresentation and the great cognitive divide: A commentary on Henriques' ' psychology defined'. *Journal of clinical psychology*, 60: 1263-1266

Stich, S. 1990. *The fragmentation of reason: Preface to a pragmatic theory of cognitive evaluation.* Cambridge: The MIT Press

Sutherland, S. 1994. *Irrationality: the enemy within.* London: Penguin

Tomasello, M. 1999. *The cultural origins of human cognition.* Cambridge: Harvard University Press

Tooby, J., Devore, I. 1987. The reconstruction of hominid behavioral evolution through strategic modelling, in Kinzey, W. (ed.). *The evolution of human behavior, primate models.* Albany: Suny Press

Tooby, J., Cosmides, L. 1992. The psychological foundations of culture, in Barkow, J, Tooby, J, Cosmides, L. (eds.) *The adapted mind: evolutionary psychology and the generation of culture.* Oxford: Oxford University Press

Toulmin, S. 1961. *Foresight and understanding: An Enquiry into the aims of science.* Bloomington: Indiana University Press

Trabant, J. 2000. How relativistic are Humboldts 'Weltansichten', in Pütz, M. & Verspoor, M. (eds.). *Explorations in linguistic relativity*. Amsterdam: John Benjamins Publishing Company

Vollmer, G. 1975. *Evolutionäre Erkenntnistheorie.* Stuggart: S. Hüzel Verlag Stuttgart

Vollmer, G. 1984. Mesocosm and objective knowledge. On problems solved by evolutionary epistemology, in: Wuketits, F. (ed.). Concepts and approaches in evolutionary epistemology. *Theory and decision library*, vol. 36. Dordrecht: D. Reidel Publishing Company.

Vollmer, G. 1987. On supposed circularities in an empirically oriented epistemology, in Radnitzky, G., Bartley, W. (eds.). *Evolutionary Epistemology, rationality, and the sociology of knowledge.* La Salle: Open Court Publishing Company

Von Humboldt, W. 1820. Ueber das vergleichende Sprachstudium in Beziehung auf die verschiedenen Epochen der Sprachentwicklung, in: Flitner, A., Giel, K. (eds.). 1963.

Wilhelm von Humboldt. Werke in fünf Bänden. Darmstadt: Wissenschaftliche Buchgesellschaft

von Uexküll, J. 1909. *Umwelt und Innenwelt der Tiere.* Berlin: J. Springer

Wallace, A. 1858. *On the tendency of varieties to depart indefinitely from the original type.* Paper on natural selection sent by Wallace to Darwin. Retrieved from: http://people.wku.edu/charles.smith/wallace/S043.htm

Warren, R. 1999. *Auditory perception: A new analysis and synthesis.* New York: Cambridge University Press

Winter, W. (ed.) 1983. *Trends in Linguistics. Studies and Monographs.* New York: Mouton Publishers

Wittgenstein, L. 1922. *Tractatus Logico-Philosophicus.* Translated by Ogden, C., Ramsey, F. London: Routledge and Kegan Paul

Whorf, B. 1956. The relation of habitual thought and behavior to language, in Carroll, J. (ed.). 1956. *Language, Thought and Reality.* Cambridge: MIT press

Wright, C. 1992. *Truth and objectivity.* Cambridge: Harvard University Press

Wuketits, F. (ed.). 1984. Concepts and approaches in evolutionary epistemology. *Theory and decision library*, vol. 36. Dordrecht: D. Reidel Publishing Company.

Wuketits, F. 1990. *Evolutionary epistemology and its implications for humankind.* Albany: State University of New York Press