Title of Ph.D. thesis:

Epistemological Prospects of Evolutionary Models of the Growth of Knowledge

Name of candidate:

Davide Vecchi

The London School of Economics and Political Science
Department of Philosophy, Logic and Scientific Method
Ph.D. in Philosophy

Abstract

In the thesis I will argue that some models of evolutionary epistemology provide an extremely illuminating and original explanation of the workings of the scientific process. Evolutionary approaches to the growth of scientific knowledge have been criticised because of the putative existence of fundamental disanalogies between biological and scientific selective processes. I will show that these criticisms are largely misguided.

I will distinguish two main kinds of evolutionary models. EEM models, which focus on the evolution of human cognitive mechanisms by natural selection (e.g. that developed by Ruse), do not provide a satisfactory basis on which to explain the nature of scientific selection processes, which are cultural rather than biological in origin. EET models, by contrast, focusing on the cultural and social origins of the selective systems operating in science, are better suited to this task. I will focus mainly on the EET models proposed by Donald Campbell and David Hull. Two general themes emerge from their analysis: the emphasis on the general validity of the variation-selection model of knowledge acquisition (i.e. trial-and-error), and the view that science is a socially adaptive and adapted system, governed by the action of peculiar selective mechanisms that partially lead to epistemic success.

On the basis of the critical examination of these EET models I will argue for three main conclusions. First, EET approaches are correct in rejecting the methodological individualism so central to many alternative epistemologies. Second, EET models offer us genuinely normative epistemological insights, particularly where social epistemology is concerned. Third, EET provides a viable naturalistic alternative to social constructivism, by justifying epistemic standards as "evolutionary constructions" (i.e., products of selection processes).

Table of Contents

Chapter 1: Evolutionary Epistemology as Naturalised Epistemology.

" Were we to rely on our current beliefs about nature in justifying the procedures of reasoning through which we arrived at those beliefs, there is a serious danger that the entire enterprise would be infected with error. Perhaps we are merely engaged in self-congratulation, when, all the while, faulty methods are being validated by the flawed conclusions to which they give rise. Naturalists have to insist that this is a genuine possibility, which cannot be excluded by invoking some set of *a priori* principles and rules of inference that are beyond criticism. We should know, in advance of skeptical embarrassments, that some forms of the problem of classical justification are solvable and others are not."
Philip Kitcher "The Advancement of Science" (p.298)

The history of epistemology of the twentieth century shows how the views of logical empiricism have been diluted, if not downright overturned. Attempts to salvage the legacy of the "received view" (cf. F. Suppe 1974) still have some prominence in some epistemological quarters.[1] However, it seems to me that the epistemological systems devised by the defenders of the "received view" are either epicyclical, or simply lack the complete set of necessary resources to solve the vast array of epistemological problems we face. In the first instance, many of these systems are based on the postulation of some implausible auxiliary assumptions concerning the nature of rationality, of epistemic standards, and the notion of scientific progress. On the other, they seem to be incapable of defeating relativistic and anti-science tendencies. What we gain is a largely unrealistic view of the practice of science, with, I believe, a consequential difficulty to fulfil the normative task so central to epistemological investigation. In order to avert such negative consequences, in this thesis I propose to look at science from the evolutionary perspective.

Evolutionary epistemologies are naturalistic epistemologies. The basic challenge for naturalistic (and therefore also evolutionary) epistemologies is, according to the defenders of the "received view", to show that their approaches to the epistemological enquiry retain the necessary normative element. I believe that the challenge must be clarified at the outset, since

[1] For instance, E. Sober (1999c) tries to show that some central ideas of the "received view" should be retained.

"epistemology" has different meanings in different idiolects. For this reason epistemologists vary in attributing to a certain problem a prominence that others regard as unjustified. Epistemology has a descriptive, an explanatory and a prescriptive dimension. No single epistemology has the resources to provide a solution or to treat the multifarious nature of the epistemological problems that need to be addressed along these three dimensions. For this reason, in this dissertation I shall advocate a pluralistic approach.

Epistemological pluralism should not be considered per se as a panacea to save epistemology from its traditional pitfalls. Many philosophers and students of science consider epistemology as a "dead" discipline. In what follows I shall try to show that this is not so. In doing so, I shall endorse the naturalistic turn. However, I am not going to argue that traditional approaches to epistemology are completely fallacious, nor that they are incompatible with the evolutionary stance. My project is humbler. The thesis I will defend is that evolutionary epistemology has much to teach us about the nature of science and that it offers many original insights that help us to explain the nature of the scientific process.

Evolutionary epistemology is a naturalistic endeavour that describes, explains and provides normative insights about the evolution of science. In this chapter I shall start the analysis of evolutionary epistemologies indirectly by focusing on epistemic naturalism. Historically, epistemic naturalism came back into favour in the 1960s. Quine (1969) has been one of its major advocates. He suggested that since Carnap's programme (i.e. the attempt to translate theoretical terms into the vocabulary of the observational basis) is unachievable, then the last reason for preferring rational reconstruction to psychology vanishes. Hence Quine's dictum (Quine 1969, p. 78): why not settle for the real thing, for psychology, and studying the relation between sensory experience and theoretical knowledge? In what follows I shall show that Quine's characterisation of epistemic naturalism is no longer up to date.

1.1 A characterisation of epistemic naturalism

Epistemic naturalism can be generally viewed as a variety of philosophical naturalism, where the latter rejects all forms of supernaturalism and transcendentalism, holding that reality, including human life and culture, is exhausted by what exists in the causal order of nature (Giere 1985).[2] Analogously, epistemic naturalism states that all forms of human inference (including the principles of logic) are natural products. As a consequence, the quest for methodological principles, principles regarding scientific inference and epistemic norms can only proceed descriptively and empirically, that is, scientifically. The idea of the continuity between science and epistemology is at the heart of epistemic naturalism. In fact, all epistemic naturalists share the continuity tenet. From an historical point of view the basic motivation for the endorsement of this tenet is that naturalists start from the acknowledgement that the Cartesian epistemological project cannot be accomplished because the absolute validation of scientific beliefs cannot be obtained. This means that epistemology cannot be conducted without presuppositions. The only open alternative is to justify them derivatively by reference to the empirically ascertained, actual-world relative and contingent reliability of our cognitive processes.[3] The naturalistic justification programme must be fulfilled by assuming the approximate truth, but yet revisable nature, of our best to date scientific theories about the nature of human cognisers and of the environment in which the cognitive ends are pursued. In the light of such contingent and factual knowledge, we decide what strategies of research and cognitive processes are likely to be reliable in producing a better picture of reality. As Donald T. Campbell (1977) puts it, the basic naturalistic question is: in a world with the naturalistic features we believe to know, and with cognisers with such biological, psychological and social characteristics, what are the best (i.e. more reliable) cognitive strategies to acquire knowledge?

---

[2] An important point is that while materialism, physicalism and other strong forms of reductionism imply naturalism, naturalism does not entail per se any of these ontological theses.

[3] This naturalistic move, coupled with the rejection of the Cartesian programme, has been heavily criticised. For instance, as Dretske (1971) points out, Darwin does not help us with Descartes. This is because, according to Dretske, epistemology is about the right to be sure, certainty and "If this sounds silly, it is because epistemology is silly....". The two alternatives, according to Dretske, should be either to give up a silly enterprise, or to try to solve the Cartesian problem. Naturalists disagree with Dretske. They do not accept that refusing to play the Cartesian game equates to the "death" of epistemology.

The continuity thesis at the heart of epistemic naturalism can be interpreted in many ways and naturalists strongly disagree about the nature of the relationship between epistemology and science. While some propose to rely on empirical evidence as the sole criterion to solve epistemological problems, others believe that conceptual analysis retains a major role to play, especially concerning general issues pertaining to meta-epistemology. In three areas in particular I see opinions diverging drastically. First, on the metaphysical issues concerned with the nature of the epistemic and its relationship to the descriptive. Second, on the issue of the roles that empirical and *a priori* knowledge have on epistemology. Third, on the issue of the naturalistic basis or (to use a more tendentious term) "foundation" of epistemology, which is also related to the issue of the methodological continuity between sciences and epistemology. I shall now consider these issues in turn.

Starting with the first, we can illustrate three ways to construct metaphysical continuity. The first is through eliminativism, where reference to the epistemic is eliminated and completely replaced by reference to the descriptive. Eliminativism sees naturalistic epistemology as a successor discipline to traditional epistemology, in the sense that the new purely descriptive epistemology does not address the same questions, since the problems of traditional epistemology are either irrelevant or unanswerable. Eliminativism implies the surrender of the normative project.

To see why, let us consider, for instance, Quine's (1969) eliminativist approach. Quine claims that, given the failure of foundationalist approaches in all forms, the epistemologist should study via psychology how theory relates to evidence. As Kim (1988) points out, what is new in Quine's proposal is not the abandonment of the quest for certainty, but, rather, the discrediting of the very conception of normative epistemology, of the justificatory project of validating scientific knowledge. The psychological theory of cognition that Quine seeks is thus better seen as a successor discipline to traditional normative epistemology.[4] A presumably

[4] Similar ideas have been proposed by other critics of traditional epistemology. For instance, Barnes and Bloor (1982) start from the analysis of the pitfalls of the Cartesian programme in its different guises and then try to revise

good reason to give up normativity might be that the naturalistic criteria of justification might happen to be highly contextual and contingent. The eliminativists are certainly correct in pointing out that assuming that a homogeneous theory of science can be achieved is a mistake. They are also correct in highlighting that science might turn out to be far too complex a phenomenon for useful generalisations about it to be achieved, and that some scientific episodes will have no deeper causal explanation than the reference to such contingencies. However, I believe their general argument is not conclusive because the success of alternative ways to construct the normative project must be judged a posteriori. I agree with Kim (1988 p. 391) when he points out that: "...for epistemology to go out of the business of justification is for it to go out of business."

Non-naturalists and naturalists without eliminativist leanings can agree that the criteria of epistemic evaluation must be grounded somehow in descriptive terms without agreeing on what kind of descriptive terms are to be employed. The problem is thus to find a way to characterise epistemic naturalism without resorting to metaphysical eliminativism.

One way is via supervenience. Kim (1988) argues that it is possible to reject epistemological naturalism in its eliminativist form without denying an almost obvious supervenience thesis. This thesis states that there are naturalistic criteria of justified belief because epistemic properties, and more generally valuational or normative ones, supervene on natural properties. The supervenience-based thesis of metaphysical continuity makes the epistemic dependent on the descriptive in the sense that the former supervenes on the latter, even though epistemic properties are neither identical to nor constituted by descriptive ones. They are fixed and metaphysically determined by naturalistic ones, but at the same time they remain *essentially* normative. This means that any attempt of reduction will fail since epistemic properties contain an essential epistemic component. In brief, two beliefs cannot be naturalistically identical and differ in epistemological value (objects indiscernible in regard to fact must be so in regard to value), and justification cannot be a brute fact of the matter unrelated to the kind of belief in question: "if a belief is justified, that must be so because it has certain factual non-

the epistemological agenda by rejecting any possibility to save the normative endeavour. What remains to be done is epistemic properties, such as perhaps that it is 'indubitable', that it is seen to be entailed by another belief that is independently justified, that it is appropriately caused by perceptual experience, or whatever.....Something like this, I think, is what we believe."[5] Supervenience allows constructing the metaphysical continuity thesis while at the same time avoiding the pitfalls of eliminativism. Supervenient naturalism can carry out a normative project even though epistemic properties are dependent on descriptive ones (e.g. explicable in terms of reliability, successful prediction, conceptual fitness).

There is another way to construct the continuity thesis: reductionism, where the epistemic is reduced to the descriptive. In this case epistemic value is a species of instrumental value, that is, a descriptive fact about a particular species of means-ends relationships. Epistemic "oughts" are identical with descriptive facts about instrumentally appropriate behaviours relative to epistemic ends. However, epistemic properties do not cease to exist after reduction (just as water does not cease to exist after being reduced to atoms of hydrogen and oxygen), and retain an unavoidable epistemological function (Maffie 1990b). The most important outcome of the reductionist construal is that epistemology becomes essentially descriptive and only hypothetically normative. Epistemology plays and should retain a normative role in virtue of its instrumental utility relative to the satisfaction of our cognitive ends. It becomes normative only within the framework of instrumental reason, in the sense that epistemic claims can be evaluated only in reference to some cognitive end. Epistemic rules are conditional imperatives rather than categorical ones. They are empirically defeasible imperatives because they are contingent claims about optimal ways to realise chosen cognitive goals. In a nutshell, normativity is grounded in instrumental reason, in facts about ourselves, our environment, our contingent ends, and in facts about what we must do to realise our ends in such an environment given our cognitive possibilities.

Passing to the second issue of the role of *a priori* epistemology, opinions diverge as well. Following Maffie's (1990a) terminology, we can thus illustrate the nature of the debate. "Limited" naturalists (e.g. Goldman 1986 p.9) deny that *a posteriori* considerations are

---

merely to analyse the scientific process in detail.

relevant to foundational or meta-epistemological questions (e.g. the determination of the proper ends of inquiry, the nature of the basic rules of justification or core epistemic values, the metaphysical status of the epistemic). On the other hand, for "unlimited" naturalists even meta-epistemology is a scientific enterprise and not an autonomous field of enquiry (cf. Laudan's rejection of "armchair epistemology"). Such a view is historically understandable as a reaction against the attempt to frame epistemology on *a priori* methods of dubious value (e.g. linguistic analysis, ordinary language philosophising, intuitionism).[6]

On the third issue, opinions diverge since different naturalists believe that some particular branch of empirical science must be granted a more basic role both pertaining the methods used to acquire knowledge and the evidential basis on which to base epistemic claims. In this light, psychology, evolutionary theory, sociology and history have all been granted this kind of role. In this thesis I shall consider evolutionary epistemologies of various forms, that all share in one way or another a commitment to treat Darwinian theory (more in particular the theory of adaptation via natural selection) as a privileged basis (cf. section 1.3).

I believe that a coherent naturalistic approach should favour metaphysical reductionism (over supervenience) and unlimited naturalism (over limited naturalism).

My reason for preferring reductionism to supervenience is the following. We can characterise epistemic naturalism either as an ontological or as a methodological thesis. In the first sense naturalism endorses a kind of ontological monism (e.g. physicalism) in the sense that only descriptive fact is assumed to exist. As a methodological thesis the idea is that only scientific method (however defined) is cognitively fruitful. In this sense naturalism should tolerate ontological diversity as long as the existing objects are accessible via scientific methods and are causally efficacious, or, in brief, integrated within the natural order. Epistemic properties must be, for naturalists, neither methodologically sui generis (requiring special cognitive methods to be grasped, i.e. non scientific) nor metaphysically sui generis (autonomous). Now, as Maffie

---

[5] Kim 1988 p. 399

[6] A clear unlimited naturalist is Rosenberg (1996), who believes that even meta-methodology has factual content. Consider also Stich's (1990) attack against analytic epistemology (cf. note 8).

(1990b pp. 288-9) points out, the problem with supervenience is that according to its theorists (e.g. Kim and Goldman) epistemic properties are conceived as possessing a sui generis normative residue that is in principle non accessible via scientific methods. For instance, Kim (1988 p. 400) thinks that epistemic properties possess a normative dimension that at some level is only accessible intuitively. More specifically, some supervenience theorists believe that the meta-level of epistemological enquiry into the nature of the epistemic is methodologically discontinuous with science, and that it requires special *a priori* non-natural methods. This means that the metaphysical nature of the epistemic is not entirely captured by naturalistic means, and, a fortiori, that there exists a mysteriously autonomous epistemic realm.[7] Hence supervenience could be inconsistent with a coherent endorsement of the thesis of the continuity between science and epistemology. Having said this, it must be stressed that reductionism is not the only naturalistically sound choice vis a vis the metaphysical question. While most naturalists endorse a nominalist and reductionist approach, agnosticism is popular (e.g. Laudan 1996 p. 165).

I favour unlimited naturalism for the same reason: thinking that meta-methodology must or could be sought via non-scientific methods (e.g. intuition) equates to rejecting the thesis of the continuity between science and epistemology at the heart of epistemic naturalism. In order to illustrate the problem with the limited approach let us consider Goldman's epistemology (1986 ch. 1). Goldman relies on a criterion of rightness in order to justify its justificatory rules. The problem is that at some point we need to decide what criterion is correct. In order to do so Goldman resorts to *a priori* meta-methodology (i.e. conceptual analysis) claiming that a criterion is to be embraced if it is coherent with our pre-theoretic intuitions about justification. That is, to stop the regress Goldman relies, like his anti-naturalist colleagues, on

[7] A further reason to reject supervenience is given by the similarity between such thesis and the realism of some critics of naturalism. Non-naturalistic approaches sometimes reject naturalism's nominalism and reductionism concerning the epistemic. An important example of this attitude is Worrall (1990) who believes in the existence of real, though non-physical, epistemic facts not reducible to descriptive ones which causally act on the epistemologist's mind in some mysterious non-physical way; epistemic facts (logical and methodological) are grasped intuitively, in a process that is akin and on a par with sense perception of physical facts; that is, epistemic evaluations are irreducible epistemic entities. Naturalists should reject this view.

intuition. I believe that Goldman's approach is not genuinely naturalistic.[8] With this I do not want to claim that conceptual analysis becomes an invalid method to acquire epistemic knowledge. It remains a useful method with the proviso that its results could only be tentatively accepted, respecting the basic tenet of epistemic naturalism, namely that justification rules must remain defeasible and corrigible independently of the ways in which they were originally acquired, and that they are to be treated as conjectures about means-ends relationships that are to be assessed empirically in terms of their instrumental utility in promoting epistemic ends. In a nutshell, epistemic theories are scientific hypotheses that are evaluated by means of the feedback relationships (which goes both ways) between norms and substantive claims. In practice, we assess norms in light of practical results, and practical results in light of norms (this is a typical process of reflective equilibrium).

The endorsement of a particular version of the continuity thesis is not sufficient to provide a sound characterisation of epistemic naturalism. The issue to be explored at this juncture is whether it is possible to identify a more substantial set of commitments shared by all naturalistic approaches. In order to pursue this investigation I shall now briefly consider some attempts to characterise epistemic naturalism that have had much resonance in the recent philosophy of science literature.

One of the most interesting characterisations of epistemic naturalism is Kitcher's (1992). Kitcher labels his naturalistic approach "traditional naturalism", since his main aim is to save part of the epistemic tradition through a naturalistic metamorphosis. The basic tenet of Kitcher's naturalism is the rejection of the *a priori*. The *a priori* analysis of the concepts of

[8] In this sense I endorse Stich's (1990 ch. 4) criticism of the limited approach to epistemology epitomised by Kim and Goldman. Strictly speaking, Stich is not interested to defend unlimited naturalism, but is rather trying to undermine Goldman's limited naturalism as a variety of the much broader analytic approach in epistemology. Stich argues that the entire edifice of analytic epistemology (which bases its account of justification on the analyses or explications of our common-sense and intuitive epistemic notions) is on shaky foundations. Stich's basic argument is that Goldman's epistemology is vitiated by a vicious circularity. According to Stich, the basic defect of Goldman's reliance on intuition is that our commonsense evaluative concepts, embedded as they are in everyday language and thought, are culturally acquired, and that therefore we have no reason to believe that the locally prevailing notions are superior to alternative intuitions. Vicious kinds of relativism would therefore be inescapable. It is important to

justification and rationality, with its appeals to conceptual truths and intuition, is banned (Kitcher 1992 pp. 11 and 16-7). The rejection of the *a priori* implies that in order to engage with epistemological research some scientific beliefs must be assumed (ibid. p. 39). This commitment to some kind of presuppositions is heavily criticised because of its consequential circularity (cf. next section).

In my opinion, the most interesting tenet of Kitcher's (1992 pp. 12-3) characterisation is the emphasis on what he calls the "meliorative project", that is, on the attempt to improve our cognitive performance in the actual world. On the one hand, this move determines an enlargement of the agenda of epistemological enquiry, since what needs to be done is not solely to show that some methods of investigation are reliable in the abstract, but more to the point that they are more reliable than other existing ones. On the other, such emphasis implies a partial rejection of the traditional task of searching for universal epistemic recommendations rather than for local ones. The important point Kitcher makes is that searching merely for universal and abstract norms limits and vitiates from the start the possibility to pursue the meliorative project.

It is interesting to note that naturalists share with many anti-naturalists the commitment to fallibilism and the view that human cognitive systems are fallible. Where they differ is in drawing the epistemic implications: while for many anti-naturalists attention to the logic of science[9] (and the related search for abstract and universal models of rationality) remains all that is epistemically relevant, for naturalists accepting the fallibilist thesis means recasting the epistemic game by developing in new directions the epistemological agenda. In this sense, the central problem of traditional naturalism becomes the maximisation of epistemic utility for cognitively limited creatures in the actual world (Kitcher 1992 p.24). In particular, naturalists refer to our limited psychological capacities and our biological heritage in order to fulfil their tasks, since epistemological prescriptions should be grounded in facts about how systems like ourselves can attain our epistemic goals in such world as ours (ibid. p.11).

---

stress that Stich does not criticise (as many anti-naturalists would do) Goldman because the latter's approach is ultimately relativistic. For Stich, who is an unlimited naturalist, relativism is true and not necessarily dangerous.

As a consequence, psychology and biology become epistemic "foundations". Psychology re-enters the epistemic game at least in the sense that we rely on psychological evidence to revise our epistemic agenda: if psychological research tells us that something cannot be the case (e.g. that humans do not behave as Bayesian agents), then our cognitive practice cannot be sought along such lines (ought not implies cannot). Biology enters the epistemic arena for parallel reasons and for an additional one: for many naturalists Darwinism at least provides a prima facie answer to the radical sceptic who claims that our epistemic edifice is rotten at its foundations. In this vein, natural selection substitutes Descartes' God in a framework that seeks to partially justify knowledge by attributing to it at least an epistemically limited survival value. In the next chapter I will show how this legitimate move to answer the sceptic can be stretched too far.

The primary focus on actual cognitive performance leads to another basic tenet of traditional naturalism: the instrumental conception of rationality. According to Kitcher, the important epistemological problem is not whether our cognitive tools are "rational" in some sense or another, but whether they are instrumental towards the attainment of our cognitive ends, whatever the latter are.[10] For instance, traditional naturalism rejects as illegitimate any attempt to justify induction *a priori* (e.g. by showing it is part of our conception of rationality, Strawson style). *A priori* analyses of this kind do not exhaust the issue of naturalistic justification, which instead amounts to establishing whether our inductive methods are reliable (Kitcher 1992 pp.11-12).

Kitcher (1992 pp. 20-1) aptly emphasises that the abandonment of the Cartesian programme has a further fundamental reason: since our knowledge depends on our reliance on the epistemic authority of others, since we depend on past historical knowledge, it follows that foundational solipsism is a failure. The recognition of the social character of knowledge is another basic tenet of Kitcher's characterisation. The meliorative project includes issues

---

[9] With the term "logic of science" I here refer to the many varied formal and historical attempts to identify the kernel of scientific method of the kind the logical positivists and empiricists seek.

pertaining to social epistemology. In a nutshell, the central problem of epistemology becomes that of understanding the epistemic quality of human cognitive performance, and to specify strategies by means of which we can improve our performance, where among such strategies Kitcher includes social ones.[11]

A straightforward way to characterise epistemic naturalism is Giere's (1985). His naturalistic philosophy of science is based on three already familiar theses: the reducibility of the epistemic to the descriptive, the instrumental nature of rationality and Darwinism. The first tenet deals with a metaphysical issue. The second is more significant epistemologically, since it states that the rationality of science is a function of the instrumental character of scientific decisions. The third thesis can be thus illustrated: naturalistic philosophy of science is based on a minimal "evolutionary" approach. In a recent paper Giere (2001) proposes to characterise epistemic naturalism as a series of methodological theses. The virtue of the methodological stance consists in not requiring a transcendental justification that naturalists cannot provide without going beyond the limits of naturalism. Commitment to the method can be somewhat justified by appeals to past successes in finding naturalistic explanations.

One methodological thesis at the heart of epistemic naturalism is "naturalistic priority", which states that the availability of a naturalistic explanation renders unnecessary any non-naturalistic ones. Naturalists must follow the method according to which all aspects of the world can be given a naturalistic explanation, among which scientific explanations are mostly relevant. A scientific explanation is an explanation sanctioned by a recognised science. This sociological criterion is the best at the naturalist's disposal to demarcate science and pseudo-science. Of course, the possibility of a naturalistic explanation for everything must remain a hypothesis subject to critical examination.

Another significant thesis in Giere's characterisation is that science only needs conditional norms, where these can be justified naturalistically. In fact, relative to a certain cognitive aim,

[10] Of course, instrumental epistemologies require an account of what constitutes cognitive virtue. I will treat the issue a little later in this section.

[11] Cf. Kitcher (1993 chapter 8) where this programme is outlined in detail.

we can in principle ascertain whether a hypothesis better fits the world than another. However, this inductive method of hypotheses selection cannot be justified a priori. In this sense, Giere (2001, p. 60) believes that the quest for a foundational inductive method that can be applied with no prior general knowledge cannot be successful.

Laudan's (1987a and 1996 esp. ch.s 7 & 9) characterisation of epistemic naturalism is clearly aimed at preserving much of the traditional agenda. In fact, Laudan labels his approach "normative" naturalism. Normative naturalism is a doctrine about the status of philosophic knowledge based on the following theses: first, the theory of knowledge is continuous with theories about the natural world; second, philosophy is neither logically prior nor superior to scientific knowledge (Laudan 1996 ch.9).[12]

The basic tenet of Laudan's epistemology is instrumentalism: methodological rules are contingent statements about instrumentalities, that is, about effective means for realising some epistemic end. Epistemic claims are not, as traditionalists sometimes believed, categorical imperatives, but rather hypothetical ones, linking some form of action to some goal. As conditional imperatives epistemic claims can only be warranted insofar as they stand in appropriate means-end relation to our cognitive ends. Their assessment is thus an empirical and contingent matter depending on the nature of the world and the nature and desires of knowers (Laudan 1987b). Testing methodological claims is in principle straightforward: the conditional statement asserting the existence of a contingent relation between means and ends (e.g. if the aim is x, then you ought to do y) refers to two observable properties (i.e. doing y and realising x). Both antecedent and consequent of the claim are thus empirical statements. The testing procedure refers to another statement that has the form of a statistical law (e.g. doing y is more likely than its alternatives to promote x).

In a way the endorsement of the instrumental theory of rationality implies the abandonment of "armchair" epistemology. In fact, given that epistemic claims can be empirically evaluated,

[12] Laudan's epistemology includes two main related parts: methodology and axiology. Here I focus on the first. But the second is crucial for many reasons. First, because, according to Laudan, aims change over time and hence a theory of aims is necessary for science. Second, because the validation process involves axiology (cf. note 14).

Laudan can show that there is no need of a special (*a priori*, supra-empirical) meta-epistemology. Thus, epistemology becomes as precarious as science itself (Laudan 1996, p. 141).[13] More specifically, the foundational science in Laudan's case is history, which has a key role in providing evidence to formulate the statistical laws linking means and ends. Instead of resorting to pre-analytic intuitions and armchair philosophising (Laudan 1996 p.138), what is needed to solve epistemological debates is just reference to the historical record. We can thus empirically find out which methods used in the past have promoted our cognitive ends (Laudan 1987b).[14]

After briefly illustrating these important characterisations of epistemic naturalism we can answer the question of whether there exists a common core of tenets typifying these different approaches.

It is clear that Kitcher, Giere and Laudan start from the endorsement of the continuity thesis, which implies both the rejection of the *a priori* (in the form of the existence of a distinctive philosophical and not scientifically-based method of epistemological investigation) and the commitment to treat empirical evidence of various kinds as epistemologically relevant. Rosenberg (1996) characterises these elements by speaking of repudiation of first philosophy and scientism (i.e. conviction that the sciences have to be the guide of the epistemological enquiry). These two elements are naturally intertwined. One way to see this link is by considering the nature of the procedure of naturalistic justification: without first philosophy and *a priori* analysis, what remains at our disposal is just contingent knowledge. As Rosenberg points out, the Archimedian pivot of epistemic naturalism in all its different forms is that epistemic justification is a contingent matter crucially dependent on the historical fact that our past ways for assessing evidence have somehow been effective in producing

[13] However, Laudan (1987b) points out that conceptual and philosophical analyses (and more generally theoretical considerations) have a role to play in assessing standards, as in physics theoretical considerations are important to assess physical claims. Normative naturalism is thus a view according to which epistemology is a mixed empirical/conceptual discipline, like the theoretical sciences.

[14] The nature of the validation process Laudan devises is reticulated: methods pick up our theories and such theories (which tell us what the world is like) tell us what sorts of methods are likely to be successful. This process of mutual adjustments also includes axiology.

knowledge (Rosenberg 1996 p. 2). The crucial point is that whether a cognitive strategy promotes a certain aim is a contingent matter dependent on how the world is.

Another linked and important element that unifies these three approaches is the commitment to an instrumental theory of rationality. The most obvious outcome of this view is that epistemic claims become conditional rather than categorical imperatives. Naturalistic normativity is only hypothetical. Another consequence of epistemic instrumentalism is that axiology passes from being a completely neglected branch of epistemology to assuming a fundamental role.[15]

Apart from these general commitments the three approaches I have illustrated diverge in considerable ways. The differences between them exemplify the rich variety of the naturalistic approach. Some of these differences are philosophically more relevant and I shall now concentrate on these.

While in Kitcher's and Giere's approaches Darwinism is a central commitment, this is certainly not the case with Laudan's normative naturalism. Perhaps this disagreement simply marks a difference in emphasis rather than a real epistemological one. But the real issue is whether Darwinism should be considered a fundamental tenet of epistemic naturalism. Rosenberg (1996) argues that it should be because, first, evolutionary theory is an excellent model of scientific theorising, second, because it is highly relevant to human affairs and, third, because the theory of natural selection provides an argument for minimally justifying our knowledge, and thus answer the sceptic (cf. section 2.4). I agree with Rosenberg. In this thesis I will focus on evolutionary models of the evolution of science that endorse a much heavier approach than a minimal commitment to Darwinism, in line with the tenets of universal selection theory (cf. section 1.3 and esp. ch. 3).[16]

[15] The instrumentalist move has been criticised because epistemic value cannot be reduced to the instrumental relationship linking means and ends (Siegel 1990). But, as Laudan (1996 pp.173-179) shows, the only difference between naturalists and non-naturalists is that for the latter there is no need to engage in axiological research, given that the fixed aim of science is assumed to be truth. But this only means that non-naturalists are not actually arguing against an instrumentalist conception of rationality. In fact, by defining the epistemic good as truth they implicitly justify rules instrumentally as good means to realise the aim of enquiry, that is, truth.

[16] The central tenet of universal selection theory is that the processes of blind variation and selective retention are universal, that is, operative in all cases of knowledge acquisition. The best formulation of this idea from an epistemological point of view is found in the writings of Campbell (cf. also Cziko). An articulation of this

Another difference between the approaches regards the reintroduction of psychology to epistemology, which is a fundamental commitment of Kitcher's traditional naturalism and Giere's (1985) "cognitive" approach, but that is absent in Laudan's normative naturalism.[17] However, arguably even this difference refers to underlying differences in the epistemological agenda rather than to genuine theoretical divergences: if your problem is to show that the scientific process can be an epistemically conducive one, then arguably psychological evidence becomes less relevant than historical evidence; but if the epistemic issue you are tackling concerns whether humans are minimally rational, then psychological evidence might be more fruitful to assess. In the first case the rejection of the *a priori* is coupled with a strong emphasis on historical and sociological research, while in the second it is coupled with psychologism.

Perhaps the most significant divergences amongst naturalists concern metaphysical issues: first, on the issue of the relation between fact and value; second, on the commitment to realism over anti-realism.

In the first case, we can certainly say that Kitcher, Giere and Laudan are not eliminativists, given that their aim is to save the normative project. If naturalism is taken to mean, as Rosenberg (1990) argues, that fact is fundamental and value just derivative, then Laudan rejects this fundamentalist view, replying that neither fact nor value is eliminable nor reducible to its counterpart, that descriptive and normative claims are on the same epistemic footing, and that given the feedback nature of the justificational relationship between rules and theories "one is well advised to be leery about asserting the justificational primacy or priority of either member of the pair." (Laudan 1996 p. 165) Kitcher, as far as I know, has no opinion on the matter. He might favour a balanced metaphysical approach that amounts, in my opinion, to some kind of agnosticism, agnosticism that mirrors the commitment to treat epistemic naturalism as a methodological thesis rather than an ontological one. Giere seems, on the other hand, more direct in embracing reductionism, at least as a methodological hypothesis, for the

conception is Dawkins' idea of universal Darwinism (e.g. natural selection is a basic and primitive force in the evolution of the universe).

same reasons (i.e. coherent endorsement of the continuity thesis) I have highlighted above. My opinion on this issue is twofold. On the one hand naturalism does not seem to me compatible with the thesis according to which there exist epistemic facts having a sort of transcendental existence that is independent of the descriptive. Second, the metaphysical issue is largely tangential: naturalism could switch between Kitcher's or Laudan's versions of agnosticism and Giere's reductionism without clear epistemological implications.

Passing to the second issue, I think we can distinguish three main realist theses for our purposes. The first is the belief in the independent existence of the objects of scientific knowledge from the knowing subject. This is a trivial ontological thesis that Kitcher, Giere and Laudan embrace.[18] The second thesis is epistemological, concerning the quality of scientific knowledge: the belief that scientific knowledge is somehow representationally true, and additionally that our theories are approximately true, or increasing in such approximation. This second thesis is much more problematic. It is for this reason that epistemic naturalists' opinions differ extensively in this respect. For instance, Kitcher is a realist in this epistemological sense, while Laudan is not. The third thesis is axiological realism, that is, the belief that truth is the aim of enquiry. The epistemological and axiological theses of scientific realism are somehow connected.[19]

In the rest of this section I will instead focus on the weaker thesis of axiological realism: should it be a commitment of epistemic naturalism? Naturalists who are committed to reliabilism (e.g. Goldman and Maffie) endorse this thesis: the aim of enquiry is the maximisation of truth and the minimisation of falsehood via the recognition of justified beliefs, which are more likely to be true than unjustified ones. The commitment to the correspondence

---

[17] According to Kitcher (1992 p 8), the crux of the post-Gettier causal accounts of justification is that the etiology of the belief is fundamental in assessing its epistemic status.

[18] Giere (2001 p. 68) tries to justify this realist commitment by adopting it as a methodological thesis. Since epistemic naturalism can only be justified naturalistically, that is, without appeal to a priori and transcendental arguments, no transcendental arguments can be given to justify such form of realism, but only the confirmation given by the historical fruitfulness of the method in leading to the construction of better representations of the world compared to the alternatives.

[19] It is possible to believe that truth is the aim of science while at the same time not believing that scientific knowledge is approximating to the truth. However, it seems to me an incoherent position to endorse the latter thesis while rejecting axiological realism. This asymmetry is determined by the fact that axiological realism is a weaker thesis. I shall treat the issue of whether epistemic naturalism should be committed to scientific realism in section 5.4. What is sufficient to say at this point is that evolutionary epistemologists, as all other naturalists, disagree on this issue.

theory of truth and to consider truth as the sole and proper aim of enquiry is both at the core of reliabilism and many other naturalistic approaches (cf. Rosenberg 1996, and Campbell & Paller 1989). But naturalists with more instrumental leanings deny that this must be the case (e.g. Laudan 1987a), while others with a penchant for relativism even contend that it cannot be demonstrated that true beliefs are epistemically more valuable than false ones (Stich 1990). Their aim is to try to substitute truth with some more attainable and recognisable goal.

Laudan's case is quite interesting. He does not consider valid the assumption that there is a set of ends constitutive of science, even if this thesis is meant as a stipulative rather than as a descriptive one. Laudan is able to identify a real axiological concern because an analysis of the historical record shows that the goals of science change over time substantially, challenging the role of axiological realism as the default thesis. Laudan (1996 p.179) thus claims that normative naturalism suggests that realist aims are less than optimal. This is shown in this way: one of the corollaries of the instrumental analysis he offers is that those ends that lack appropriate means for their realisation become suspect; realist aims are highly suspect at least because there is genuine disagreement between epistemologists about which is the observational measure of truth: is it empirical adequacy, problem-solving effectiveness, falsifiability, predictive reliability, etc.? Laudan then continues that if the epistemic principle *cannot implies ought not* is accepted then we are rationally forced to abandon a realist axiology since true knowledge is not provably achievable. The outcome is that truth cannot be the cognitive aim of enquiry.

In my opinion, Kitcher's axiological realism and Laudan's anti-realism are both compatible with a naturalistic approach to epistemology. In fact, there is something odd and anti-naturalistic in determining *a priori* what the aims of enquiry should be. From a naturalistic point of view, what should be done is, first, to investigate historically which specific realist and non-realist goals have been pursued by scientists, and, second, the empirical consequences of pursuing such goals in a comparative fashion. The choice of goals becomes then a matter of practical choice informed by empirical research on what can in fact be achieved (cf. section 5.4).

From this section it should be clear that it is difficult to identify a common core of tenets shared by all kinds of epistemic naturalism. However, it seems to me that naturalised epistemology should both be committed to a conception of instrumental rationality and to a revised version of the normative project. Furthermore, it should not propose to substitute the normative with the descriptive, surrendering the normative tout court to become a pure descriptive theory of science. In the next section I shall focus on the general problems of naturalism.

1.2 Epistemic naturalism and its problems

In this section I shall focus on the three most important problems affecting naturalistic approaches. The first problem concerns the circularity of the naturalistic approach: if our epistemology is derived from the contingent information about the world provided by the sciences, then circularity ensues because our empirical information is validated by the epistemology derived from it. Secondly, epistemic naturalism supposedly opens an epistemic gap by extrapolating norms from facts, with the putative consequence that its normative project is fallacious. Thirdly, naturalists base their epistemic claims on the actuality of the human perspective rather than on an abstract epistemic situation, thus risking embracing vicious kinds of relativism.

The first two problems are more aptly seen as variants of the same one, regarding the status of epistemic norms and their legitimacy: since epistemic naturalism relies on established scientific theories to capture the normative principles, then it either begs the crucial epistemological question or it is viciously circular. Relativism seems to be a corollary of the lack of a fundamental distinction between fact and value: given that naturalists cannot discriminate between described rules governing actual scientific practice and fundamentally normative criteria, then all knowledge claims are both justified equally well (e.g. none at all) and are equally scientific (or pseudo-scientific). I shall now consider these criticisms in turn.

1.2.1 Scientism or dogmatism?

The circle argument states that to use scientific methods and scientific evaluative criteria to investigate scientific practice must either beg the question or lead to infinite regress. To stop the circle, the argument goes, we must dogmatically rely on some basic and primitive criteria of evaluation that do not require further justification, thus resorting to *a priori* considerations. While to stop the regress, the argument continues, we must rely on some methods and criteria before starting the investigation, thus acknowledging that not all methods and criteria can be discovered by epistemological investigation qua scientific enquiry. There is something wrong, the argument says, with scientism and the idea of epistemology being completely a posteriori. What we need to investigate is whether naturalists can aptly justify some basic tenets of their approach, namely scientism, the general idea that the Cartesian epistemological programme has failed, and the belief that epistemology cannot be approached by eschewing reference to some kind of presupposition.

The first thing to note is that a kind of circularity affects not only naturalistic but any kind of epistemology. As Bradie remarks (1989 p. 402), following Chisholm, there are two essential ways in which an epistemology can be constructed, which arguably correspond to the two main epistemological problems. The first is the problem of the extent of knowledge: what do we know? The second is the criteriological problem: what are the criteria of knowledge? If it is assumed that there are examples of knowledge, then a criterion can be extracted. How to judge whether such a criterion is good? It is good insofar as it produces good items of knowledge. Conversely, if a criterion is assumed, then items of knowledge can be constructed. How to judge whether these items are examples of good knowledge? This can be done because the criterion sanctions their validity.

All of this means that, in constructing an epistemology, a kind of circularity seems inescapable and unavoidable. Circularity can be completely avoided only on pain of dogmatism of some kind, or by simply denying that epistemology has any sense at all. I believe the second strategy is ludicrous, while a dogmatic standpoint is incompatible with naturalism. This is why the only naturalistic open alternative to dogmatism is scientism. The real problem is then to justify the commitment to scientism.

In a way, scientism seems a rather uncontroversial thesis that does not seem to need further justification, since science provides us with some of our best examples of knowledge. Scientism in this sense is a thesis endorsed both by naturalist and anti-naturalist philosophers of science alike. In another sense, scientism is the thesis according to which the sciences have to be the guide to the epistemological enquiry (Rosenberg 1996 p.25). This is a fully-fledged naturalistic thesis that we need to scrutinise.

As I see it, there are two problems with scientism. The first regards the justification of the idea that science is a superior epistemic practice. The second concerns the details of the scientistic outlook (e.g. Darwinism), but not scientism in general.

The first criticism is especially raised by relativists. The problem regards the possibility of justifying the assumed epistemic superiority of science. In order to provide such justification there are two main options: naturalism and *a priorism*. Naturalists argue that the justification must be empirical. Rosenberg (1996 pp. 25-6) argues that naturalists should point to the impressive record of prediction and control of science vis a vis any other epistemic practice and that only by focusing on such virtues of science the social constructivist's challenge can be rebutted. *A priorists* refuse the naturalistic move. As Worrall (1999) shows, Rosenberg's attempt to refer to the formidable record of prediction and control of science is vitiated by the acceptance of certain standards of judgement that can be easily discarded by relativists, who could ask naturalists, for instance, to justify the evaluative standards on which they sanction the superior predictive power of evolutionary theory vis a vis creationism. Thus, according to Worrall (1999 p. 8) the only good defence of scientism comes from *a priori* considerations.

The divergence between naturalists and *a priorists* is not about the truth of scientism, but rather about the ways in which the epistemic superiority of science should be vindicated.

More specifically, the nature of the disagreement concerns the role that empirical evidence should play in such vindication. Can we establish empirically the superiority of new science vis a vis old science (e.g. scientific biology vis a vis creationism) as naturalists believe? Worrall argues that at some point the justificatory process must end, specifically when we reach the basic rules of the epistemic game, where no further debate is needed: either you accept the validity of the basic epistemic rules and act rationally, or you don't and act irrationally. Naturalists reject dogmatism and *a priorism.* But they still need to give an answer to this fundamental question: where does the justificatory role of empirical evidence end?

Naturalists answer this question in many ways. The *leitmotif* of all these answers is to couple the rejection of dogmatism and the refusal to play the Cartesian game of answering the sceptic with either the endorsement of a very primitive inductive strategy or with Darwinism. For instance, Laudan's solution is to stop the justificatory regress by relying on a supposedly neutral rule of enumerative induction. This rule, named R1, is the quasi-Archimedean standpoint, the only uncontroversial rule that can be identified: "Indeed, if R1 is not sound, no general rule is." (Laudan 1996 p.135) Laudan seems to think that the circle he proposes is less vicious just because all methodologists seem to be committed to R1 as the minimal and universal requirement. But obviously this move remains subject to criticism.[20]

Kitcher does not rely on enumerative but on eliminative induction to escape circularity. The basic idea behind inductive eliminativism is that the information about the past guides our future performance by regulating the selection of hypotheses, and that the logic of inductive justification lies in the elimination of all plausible alternatives. In brief, hypotheses selection consists in the hopeful and long-term elimination of all but one of the set of hypotheses that are consistent with the evidence. Inductive eliminativism is both intended as a description of

[20] On the anti-naturalist front, Worrall (1999 pp. 13-4) argues that the adoption of R1 can only be justified if we accept his point that some standards have a supra-scientific status. On the naturalist one, Rosenberg argues that Laudan's endorsement of R1 is insufficient to provide a sound naturalistic epistemology. In fact, Rosenberg continues, epistemic evaluation must rely on other criteria (coherence, simplicity, explanatory power, predictive fertility) other than R1 that an epistemology, like any other scientific theory, must meet. Circularity is avoided, Rosenberg continues, by endorsing (not a priori, but tentatively) a proper empirical and revisable theory of enquiry. I believe that Rosenberg is right in this respect (cf. section 5.2.1).

practice and normatively. How does reliance on eliminative induction break the circle of justification? One could try to show that no other method is more rational than selective elimination (Reichenbach style). Even though I am sympathetic with these abstract demonstrations, naturalists generally try to find encouragement in Darwin.[21] In this vein, Kitcher argues that the eliminative propensity (the human tendency to think in terms of elimination of rivals) must have an evolutionary origin (i.e. humans are selected for eliminative induction), and that it was certainly efficient for our ancestors in order to cope with certain environmental challenges. In an analogous vein, Giere (1985) argues that the only open naturalistic way out of the circle is Darwinism, since Darwinism at least partly justifies the contingent information on which naturalism bases its epistemic quest. Darwinism helps naturalists to dismiss the sceptic claim that our initial cognitive state is so bad as to render knowledge impossible. In fact, evolutionary theory assures us that natural selection would have wiped out creatures with extremely fallacious cognitive tendencies.[22]

In a certain way, the Darwinian move remains open to criticism. This is because, as Worrall (1999 p. 9) argues, in accepting Darwinian theory as a kind of epistemic foundation, naturalists are, logically speaking, making undefended assumptions about the evidential status of a scientific theory. We thus see that Darwinism as a specification of the scientistic framework of naturalism does not disentangle the circle.

We have seen so far that naturalists justify scientism either through the endorsement of some primitive form of inductive practice or with some kind of Darwinian argument. We have also seen that these moves are insufficient because the commitment to some kind of primitive inductive strategy is hard to justify empirically (be it along Darwinian lines or through historical evidence), while evolutionary theory cannot be empirically justified without raising

[21] In section 5.2.2 I will return to these issues. I agree that the commitment to eliminativism is not enough to secure the objectivity of scientific knowledge. I also acknowledge that eliminativism faces many challenges. However, as I shall try to show, I believe that there is no better alternative to this method.

[22] In section 5.2.2 I will try to articulate a stronger version of the Darwinian argument in order to justify inductive eliminativism It is obvious that behaving scientifically does not equate to behaving in accordance to the eliminative rule of induction; something more is needed to be a good scientist. It is also obvious that the elimination of alternative hypotheses in science is not completely regulated by an adaptive and adapted cognitive character; science is a social process. Thus, I think that the only way to justify the basic rule of enumerative induction is by showing that it either approximates the optimal cognitive strategy or that it has no better alternative. I believe that evolutionary epistemology has some contribution to make in these respects.

a problem of circularity. We can therefore conclude that even though scientism is not a problem peculiar to epistemic naturalism it nonetheless leads to unavoidable forms of circularity. Naturalists must accept that this could not be otherwise.[23]

So the issue becomes: is there any alternative to the dogmatism and *a priorism* Worrall favours, which risks contaminating the unlimited form of naturalism I have characterised so far? In order to consider this issue let us briefly characterise the dogmatic approach. Worrall criticises the naturalistic approach because he thinks that not all epistemic standards are evidence dependent and revisable through contingent information. In this sense, he argues, naturalists fail to distinguish two basic kinds of standards: those that change (e.g. avoid hypotheses that posit unobservable entities) and those, implicit and primitive, that do not (e.g. *modus ponens*). The basic and fundamental ones must be unrevisable, in the sense that the state of the world could not determine their revision (Worrall 1999 p. 12). For these reasons Worrall favours dogmatism over naturalism.

On the naturalist front the agreement is that "dogma" is certainly too strong a word to define the nature of the theoretical tenets and working assumptions that constitute the 'foundation' of epistemology. Such tenets and assumptions must remain independently testable and revisable, at least in principle. As a matter of fact, I think the naturalist has another weapon to dispose of dogmatism. As Giere (1985 p. 336) argues, naturalists should start by questioning the presupposition that it is possible to establish and justify norms *a priori*, logically and intuitively. In this sense they should point, in order to sustain their argument, at the 300 years of failed attempts to justify inductive inference in a non-question begging way.[24]

To conclude, I believe that in part the difference between the naturalist and the dogmatist is only of attitude: some beliefs crucial to our picture of the world can either be seen as

[23] As Giere (2001, p. 60) points out, inductive foundationalism is not a viable naturalistic reply to the objection of the infinite regress of the naturalistic process of justification. This is because no inductive method can be applied without prior knowledge and thus justified a priori.
[24] The problem of induction is a puzzle that continues to fascinate philosophers. Some contemporary philosophers believe that they have solved the Humean puzzle (e.g. Howson 2001).

nonnegotiable or as revisable posits. I prefer not to overemphasise the role of this pragmatic difference, which mirrors a metaphysical debate that cannot be easily adjudicated. Again, we see that what is at stake is the belief in the continuity between science and epistemology, and in particular the metaphysical idea that facts and values belong to the natural world. After all, Worrall (1999 p. 10) acknowledges that the only (significant?) difference between the circular (i.e. naturalistic) and the dogmatic approach to defend the epistemic specialness of science is of openness and honesty. I agree with Worrall with one caveat: intellectual honesty means for me acknowledging that some standards of evaluation that we at the present time deem unrevisable might in the future be changed. What is more intellectually "honest" than tentatively embracing a perspective consisting of our best theories and methods of enquiry, by supposing their truth while retaining their fallibility, and believing that the only possible justification of science can come from "within" science?

1.2.2 Naturalistic normativity

The second and biggest challenge to the naturalistic approach is peculiar to it: since epistemic naturalism is a description of scientific practice, how can it provide and justify norms? Epistemic naturalism seems to imply either the surrender of the traditional normative project or committing the naturalistic fallacy of deriving norms from facts.

The traditional normative project was based on the possibility of extrapolating epistemic norms in the form of unconditional categorical imperatives. Naturalists carry out a significant revision of the normative project along instrumentalist lines (i.e. hypothetical normativity) but retain the classical project of seeking norms. Naturalism starts from the analysis of scientific practice. The preliminary aim is to produce adequate explanations and fruitful generalisations about how science evolves. The ultimate aim is to derive normative recommendations from the empirical claims thus obtained. Naturalistic epistemologies have therefore a descriptive, an explanatory and a normative dimension. The nature of the explanatory generalisations and normative recommendations sought by naturalists varies extensively.

The possibility of reframing in a naturalistic form the normative project is open to two basic kinds of criticism. First, anti-naturalists argue that such revision either amounts to a total surrender of the normative (because the fundamental problems of epistemology remain untreated by naturalism – cf. e.g. Kim 1988, Dretske 1971), or that it does not yield any interesting normative recommendations. Second, radical naturalists (i.e. eliminativists) deny any normative role to empirical studies of science (e.g. Barnes and Bloor 1982). I shall now focus on the particular way in which the kind of epistemic naturalism I characterised in the previous section retains a normative aspect.

The first feature of normative naturalism is the emphasis on what Kitcher calls the meliorative project. In this sense, what naturalists try to provide is a descriptive and evaluative account of human cognitive performance with the ultimate aim of improving such performance. On the one hand, this certainly means that the project is more humble than what anti-naturalists look for, but, on the other hand, it is also arguably more relevant to human affairs. As Giere (2000) says: "The general goal is not to ground what we think we know, but to improve upon it."
The revision in this sense of the normative project is coupled with the abandonment of the logic of science. The starting point of the naturalist enterprise is that, despite accepting that through logic and probability theory we can formulate cognitive strategies that are cognitively optimal, an epistemology of the cognitively optimal has many limits. First, the certainly valuable epistemic standards formulated through logic are insufficient to explain scientific decision-making.[25] Second, the logic of science has limited instrumental value. The search for an optimal account of cognition might turn out to be less relevant than a more localised and relativistic account of cognitive performance. In fact, what matters are cognitively realisable aims instead of putatively unrealisable optimal ones. From our limited human perspective, the real epistemological problem is not the achievement of an unfeasible optimality in cognitive performance but rather its improvement.

[25] As already highlighted by Kuhn, who emphasised the ideological (e.g. social) rather than the algorithmic (e.g. logical) nature of methodology.

The abandonment of the traditional normative project (with its unsubstantiated assumption about the explanatory and normative sufficiency of the logic of science, and its emphasis on the search for cognitively optimal standards) determines the naturalist's move towards the sciences. Anti-naturalists typically eschew reference to the empirical sciences because for them epistemology should concern the psychology of the "ideal mind", that is, with uncovering the laws of rational scientific behaviour (Siegel 1980).[26] Naturalists believe that the anti-naturalist reconstruction of the scientific process is partly fictional, epicyclical and ad-hoc. The only open alternative is to acknowledge that the sciences provide us with the best tool at our disposal to ascertain whether a cognitive process is aim-conducive. Thus, the naturalist's aim is to assess empirically whether the cognitive processes governing scientific practice are conducive and how to improve their efficacy. It is for this reason that empirical information about the actual world has normative relevance (Kitcher 1992 p. 35). In particular, psychology, biology, sociology and history assume an important epistemological role.

The second feature of normative naturalism is to encompass the normative force of epistemology within the framework of instrumental reason. In this sense, the naturalistic central move to save the normative project (central in the approaches of Laudan as well as Giere and Kitcher) is to ground the normativity of epistemic claims on their instrumental utility together with contingent facts about ourselves, our place in nature and our cognitive ends. The normativity of epistemology becomes contingent, hypothetical and parasitic upon that of instrumental reason. The interesting naturalistic question is: does following a certain norm promote some specified epistemic end better than its alternatives? Epistemological investigation should aim at formulating useful generalisations concerning aim-conducive scientific behaviour. The ultimate aim of the naturalistic theory of science is to provide normative guidance for scientists and for science-policy makers, fulfilling its basic meliorative task. The emphasis on the improvement of our knowledge is also shown by a significant

[26] The idea behind this thesis is that scientists act rationally only when they behave according to the dictates of the "ideal mind", that is, when their behaviour is "unclouded" by psychological or sociological factors. Only "good" scientists in their "best" moments have a sound "intuitive grasp" of methodological facts (cf. J. Worrall – 1990 pp.316-7).

corollary of the thesis: our view of what is an efficient mean to attain a certain aim will change over time and will hopefully become more accurate.

It is certainly true that naturalists have very liberal ideas about what epistemology should be about and that they pursue a peculiar epistemological agenda. The emphasis on human knowledge and not on strictly philosophical questions is at the core of their agenda. Naturalists and traditionalists put different emphasis on different sets of questions and it might be argued that in many ways the two research programmes are complementary. For many critics this is not, however, the case. Two basic criticisms of normative naturalism can be raised.

The first criticism can be illustrated as follows. We have seen that hypothetical normativity is the mark of epistemic naturalism. This position is coupled with the additional thesis stating that scientific epistemology does not require any categorical norms. Giere (2001 p. 58) thus summarises this basic point: it is true that within a naturalistic framework one cannot justify categorical norms, but only explain them. However, he continues, "conditional norms......can be justified naturalistically, and science requires only conditional norms". This is a contentious thesis that has been forcefully criticised by anti-naturalists. For instance, Worrall contends that naturalists, by failing to distinguish two kinds of epistemic standards (on the one hand categorical, *a priori* and supra-empirical ones, on the other conditional, a posteriori and open to revision ones), are as a consequence unable to sustain a vigorous normative approach that can be used to properly rebut the relativistic challenge. But naturalists should argue that any evaluative principle is at least in part open to revision, and deny that the postulated fundamental distinction between procedural and substantive standards corresponds to reality. In fact, as Laudan (1996 pp.171-2) shows, all principles of evaluation make some substantive assumptions about the structure of the world we live in and about us as thinking beings.

The second criticism deals with the naturalist conviction that descriptive knowledge of many kinds potentially helps us in improving human cognitive performance. However, to derive norms from facts means committing the naturalistic fallacy. Anti-naturalists invariably refuse to conceive norms as fallible posits or conjectures, open to empirical assessment in the light of new evidence like theories. In brief, what they contend is that descriptive and prescriptive claims cannot be subjected to the same kind of adjudication since they belong to two different worlds. To this criticism, so strongly based on a metaphysical distinction, naturalists can either refuse to see a problem,[27] or try to (dis)solve it.[28]

Naturalists agree that the sciences provide a more feasible and alternative way to ascertain whether a cognitive process is aim-conducive. But they disagree on three fronts. First, on what they understand by "cognitive process". Second, on axiological issues. Third, about the nature of the norms to be sought and the empirical evidence relevant to establish the status of norms.

On the first front, some naturalists refer exclusively to psychological processes. Others prefer a much more extended approach, including, for instance, apart from individually instantiated processes, also social ones. In the latter sense epistemic naturalism engages with sciences other than psychology (e.g. sociology). Kitcher's traditional naturalism is an important illustration of the extended approach. He argues that as soon as we accept that we rely on our predecessors and that Crusonian epistemology is just an idealisation, then we must investigate the ways in which the primary locus of scientific evolution, namely the community of scientists, advances its epistemic ends. He thus calls for a modification of the normative project along social epistemological lines (Kitcher 1992 p.61). The task of social epistemology is to study the general patterns of the emergence of genuine consensus in scientific debates, to identify the kinds of social factors that contribute to or detract from the attainment of the epistemic good, to study the role of authority and trust, the nature of cooperation and competition in science, and the role of cognitive variation in science. Furthermore, social epistemology should provide a characterisation of a well-functioning cognitive community.[29]

[27] As Laudan (1996 p. 156) puts it: "where's the fun in being a naturalist, if one is not thereby licensed to commit the fallacy?"
[28] In section 5.2.1 I shall show that the naturalistic fallacy is not a real threat to epistemic naturalism.
[29] Campbell's and Hull's evolutionary epistemologies pursue such programme in detail as I shall show in chapters 3, 4 and 5.

On the second issue, opinions diverge because naturalism is not per se committed to axiological realism. Even though there are many similarities between Giere's, Kitcher's and Laudan's approaches, these do not concern axiological matters. Kitcher tries to link traditional naturalism to realism, Giere's approach is pragmatic, while Laudan tries to sever the link between the two. Laudan's opinions are well known on the matter and have been already illustrated in section 1.1. The associated theses that the aims of science change over time and that, above all, truth is not the basic cognitive virtue, have been severely criticised. In particular, critics contend that Laudan's normative naturalism loses its normative force unless supplemented with a proper account of cognitive virtuosity that Laudan does not deliver.[30] In order to obviate such deficiency, naturalists could follow Giere (2001 pp. 57-60), who proposes an alternative to axiological realism. He argues that axiological realism is an oversimplified thesis because the notion of truth is epistemically redundant. Giere suggests replacing the notion of truth with the notion of fit or similarity between a representation and the world. By appealing to the notion of fit, the supposedly primary problem of identifying whether there exists a general antecedent clause of the various conditional norms sought by naturalism is eschewed.[31]

On the issue of the nature of the norms to be sought, the extension of the disagreement is a function both of the evidence that is taken to be most relevant and of the interpretation of the recommendations naturalists try to obtain. The first variability is obvious given that the scientific discipline of reference differs for every naturalist. The second variability depends on the epistemic aims naturalists try to seek. Some seek general methodological principles (e.g. Laudan), others more general patterns of normative behaviour (e.g. social epistemologists), while others (e.g. radical relativists) believe that it is impossible to extrapolate historically general principles, given that scientific practice reflects highly idiosyncratic patterns of behaviour.

[30] Laudan has been criticised by both naturalists and not. Worral (1999 p. 13) argues that without restrictions on aims Laudan's naturalism loses in normativity, while by acknowledging such restrictions it loses in naturalisation. While Rosenberg (1990), a strong naturalist, criticises Laudan because the suppressed antecedent clause of any conditional and hypothetical epistemic claim is inevitably the attainment of true knowedge. For Rosenberg (at least in this paper) no higher-level axiological debate is needed.

[31] This view resembles in some respects Campbell's evolutionary approach to the subject (cf. section 5.3 and 5.4).

1.2.3 Normative cognitive pluralism

The third main problem of epistemic naturalism is relativism. I believe that some form of relativism is an inevitable outcome of the naturalistic turn, but the genuine issue is whether this form is dangerous. I do not think it necessarily is.

The term "epistemic relativism", as Stich (2000) has pointed out, "has been used in a bewildering variety of ways". The significant sense of the term "epistemic relativism", Stich argues, should refer to the thesis according to which there is no unique set of evaluative standards that people ought to use, thesis that he aptly labels normative cognitive pluralism. So, whenever sensitivity to facts about an individual or a group of epistemic agents is displayed in using a set of standards, we can speak of epistemic relativism.[32] The alternative to epistemic relativism is normative cognitive monism, that is, a universalistic conception stating that there is a single set of epistemic standards that epistemic agents ought to use. Normative cognitive monism is a much stronger thesis than normative cognitive pluralism. It is also a monolithic thesis, while the latter conception comes in degrees, depending on how much universality is granted. The putative risks connected with relativism are that without universal standards of epistemic evaluation we cannot assess claims of epistemic practices that do not share standards in common with us, that relativistic standards might become non-applicable across cultures and contexts, and that inter-cultural methodological and theoretical disagreements cannot be rationally decidable.

It is certainly true that epistemic naturalism is relativistic in some sense, given that it is not committed to providing a universalistic account of cognitive value. In fact, some form of relativism follows from the naturalist's choice of reframing the normative project as an

[32] Kuhn was thought of as embracing a relativist view of this kind when he pointed out that the best criterion of scientific objectivity is the judgement of the group of trained scientists constituting the relevant scientific community. In fact, if judgements concerning the objectivity of scientific claims is sensitive to facts about the groups of scientists constituting the community at a given time, then relativism ensues: the universality of the standards of epistemic evaluation cannot be identified via a sociological (and for this reason relativistic) criterion. I believe that this thesis only becomes dangerously relativistic when coupled with the other relativistic idea that the justification of the epistemic standards is irremediably paradigm-dependent (cf. ff. note)

instrumental question about the improvement of our cognitive performance: given that naturalistic norms are hypothetical and dependent on contingent hypotheses about humans, then, as a matter of fact, the scientifically extrapolated epistemic principles necessary to improve cognitive performance in the actual world might turn out to be local. This remains an open question. The abandonment of the traditional normative project with its search for principles of universal applicability and the endorsement of the normative project in the naturalistic form is not deemed to fail because it is "somehow" relativistic. The genuine issues concern, first, the nature of this localism and, second, whether localism implies that epistemic naturalism is anormative.

The first problem could be articulated as follows: are weak forms of relativism dangerous? As already noted above, normative cognitive pluralism comes in many forms, some more relativistic than others. The spectrum of possibilities varies from radical relativism (i.e. the process of hypotheses selection is always arbitrary and unregulated, inter-paradigmatic methodological and theoretical disagreements cannot be rationally decidable, epistemic standards of evaluation are irremediably community-dependent and judgements concerning scientific objectivity have no value at all) to weaker forms of relativism stating that scientific change is "sometimes" underdetermined by good reasons. The issue is then whether "often" means "more often than not".

Concerning the second problem, normative naturalists should argue that even though some form of relativisation is inescapable, this does not imply that local normative epistemology is a meaningless endeavour. Rather, as Kitcher (1992 p.49) argues, naturalists should welcome this localism, which provides "a salutary counter to universalistic yearnings". What is needed is rather a demonstration that even non-universalistic accounts of cognitive virtue are sufficient to save normativity and extrapolate objective standards.

Two general arguments against naturalistic relativism can be proposed. The first concerns the notion of hypothetical normativity: treating epistemic claims as hypothetical rather than categorical imperatives means that naturalistic norms are conditional on many factors (e.g. our biology, psychology, cultural organisation). This hypothetical character detracts from their universality. The naturalistic justification of such norms will be a matter related to their aim conduciveness, which will be evaluated empirically via the accepted empirical practices and in relation to the chosen aims. The risk is thus that of having a validation process that is highly dependent on the epistemic situation of the scientific community at a certain time and place, that is, relative to a peculiar historical, spatial and even discipline-dependent perspective. In other words, the risk is that of finding ourselves constrained to what can be labelled a "paradigm-dependent" perspective. In such a situation, the critics of naturalism argue, the objectivity of science really becomes just a matter of persuading and converting members of different paradigms to another. I believe that since naturalists cannot give up hypothetical normativity, the only way to save the objectivity of science is to base the validation process on those cognitive factors that make it possible to transcend the paradigm-dependent community. Naturalists should thus concentrate on identifying these supra-paradigmatic and community-transcendent features that render science different and superior to all the other epistemic practices, and that allow scientific knowledge to be validated independently of what a certain community of trained scientists believes.[33]

The second general argument deals with the lack of universalism of the naturalistic approach, which may appear to yield vicious kinds of cultural and species-relativism. As a consequence, anti-naturalists argue that naturalistic normativity is either deficient or just an exercise in empty moralising.

In some forms the argument is not threatening. Consider first Kim's (1988) argument against Quine's eliminativist approach: if epistemology becomes, as Quine argued, psychology, then the emphasis on the nomological patterns relating theory and evidence are going to be species-relative rather than objective, with the subsequent failure to find universal criteria of epistemic value. I partly agree with Kim. I agree in the sense that I too consider Quine's psychologism

[33] As I show in section 5.3, where I identify the kinds of cognitive factors that render science superior to other practices. In that section I will also explain why I personally do not have a problem with a criterion of sociological demarcation of science from pseudo-science as that devised by Kuhn (cf. preceding note). The only caveat regards the idea of scientific community, which cannot be characterised in Kuhn's too restrictive terms: the scientific community must be structured more as an open society rather than a religious sect in order to efficiently pursue the cognitive good.

insufficient to ground an entire epistemology.[34] But I do not see a problem with species-relativism. It is a problem for Kim because he believes that the relation between theory and evidence of crucial epistemological interest regards abstract factors (i.e. objective and logical ones), and that in order to identify such a relation *a priori* meta-epistemology is necessary and sufficient. But this argument is weak because logic is insufficient to ground a sufficient number of powerful universal norms.[35] In brief, there is no alternative to innocuous species-relativism (cf. section 5.3).

Worrall (1999 p.3) proposes a more significant argument to attack the unpalatable cultural relativistic tendencies intrinsic in the naturalistic approach. Naturalists base their evaluative judgements on the sciences. One basic problem of this approach is: what counts as a science? Some naturalists seem to believe that the best criterion to demarcate is merely sociological. As Giere (2001 p. 55) puts it: "A scientific explanation is an explanation sanctioned by a recognized science. To say more is to go beyond the bounds of naturalism." But the problem remains that the authority to identify genuine scientific explanations, that is, to distinguish science from pseudo-science, lies in the hands of the practitioners. It is for this reason that Worrall contends that if we study science scientifically it follows that we cannot establish the epistemic specialness of science on purely descriptive grounds. What we can do, Worrall continues, is merely to highlight the existence of a difference in scientific practice. Thus, Worrall concludes, there is no good reason to study scientists instead of scientologists.

More in detail, Worrall argues in two respects. First he claims that no unconditional claim that some set of epistemic standards are correct can be underwritten by relying on a simply descriptive account of scientific practice, for the reason that it depends on the category of practitioners taken into descriptive account. If standards are open to empirical revision then it is purely a sociological fact of the matter whether we follow the practitioners in revising them.

---

[34] Naturalism in the epistemology of science does not equate to psychologism because not all cognitive processes operative in science are psychologically instantiated or describable in psychological terms.

[35] First, modus ponens (and related logical norms) is not sufficient to explain and regulate scientific behaviour. Furthermore, the assumption that *modus ponens* is universal in value rather than species-relative is taken for granted but should be left open to revision, as naturalism dictates. However, the difference between treating modus ponens as species-relative rather than universal is pragmatically obscure unless we engage in empirical investigations (e.g. do other primate species behave according to *modus ponens*?). Again, this is what naturalism dictates. I also doubt that a more substantial logic of science (Bayesianism) could explain the intricacies of scientific practice.

Some practitioners act in a supposedly scientific way, others in a supposedly non-scientific one, but naturalists cannot adjudicate the correctness of one method unless they rely on some non-naturalised assumption (Worrall 1999 pp. 6-7). Second, he claims that the naturalist view that all rules of evaluation are corrigible parts of our system of knowledge, that there do not exist supra-paradigmatic criteria of evaluation, which are unrevisable, *a priori* and lying outside the scientific game, has a deleterious outcome, since without constitutive principles of the "logic of evidence" naturalists cannot demarcate science from pseudo-science (ibid. pp. 5-6). The only way to escape relativism is to give up the purely descriptive undertaking by choosing to rely on some kinds of *a priori* evaluations (ibid. p. 4). Worrall's solution is to think that some standards of a formal, categorical and unrevisable nature are constitutive of epistemology (i.e. the logic of evidence). Relativism can only be defeated by referring to such a background of non-naturalised general principles (ibid. p. 17).

In a nutshell, Worrall believes that the reason for the failure of the naturalistic normative project is the endorsement of the continuity thesis with the consequential refusal to endorse a minimal and healthy *a priorism*. I understand the nature of this criticism and sympathise somehow with its motivation. If no *a priori* and primitive evaluative standards exist, if *a priori* epistemology is completely abandoned, if value is metaphysically dependent on fact, then unconditional norms cannot be sought, and therefore relativism cannot be defeated. But I do not think that naturalists would feel threatened by Worrall's argument. This is because they would not easily give up the criteria of selection Worrall refers to (e.g. *modus ponens* or the method to test hypotheses against all plausible rivals), whose reliability has been proved almost beyond any reasonable doubt in 300 years of scientific practice, and whose centrality in our system of knowledge is so well established as to be open to revision just in principle. This means that naturalists would refer to similar criteria of selection to Worrall's in order to justify their preference for scientific psychology over scientology (cf. section 5.3).

Worrall's argument can be interpreted in two ways. In one sense, the disagreement between the naturalist and the dogmatist is just rhetorical, given that many naturalists would give the same good reasons identified by Worrall to prefer science to pseudo-science. In this first respect, I believe that Worrall's criticism of naturalism is too strongly based on the rejection of the continuity thesis. I have already pointed out in section 1.1 that in my opinion the metaphysical dispute about the nature of the epistemic is, paraphrasing logical positivists, merely pseudo-philosophical: a difference in metaphysics does not necessarily amount to a genuine epistemological one.

In a second sense, I think that Worrall's argument points to a genuine difference between naturalists and anti-naturalists. Here the issue is not about being able to give non-trivial reasons to prefer science to pseudo-science, but it is rather about how to justify our reliance on certain standards of epistemic evaluation. The anti-naturalist refuses to accept that a purely sociological criterion can demarcate science from pseudo-science. Some naturalists retort that they still prefer the judgement of a well-structured and well-functioning scientific community to that of an *a priori* epistemologist. I think that a compromise between the two positions is more sensible (cf. section 5.3).

After having considered the basics of the naturalistic approach, it is now time to move to treat directly evolutionary epistemology.

1.3 The basic tenets of evolutionary epistemology

In the last two sections I tried to introduce indirectly the basic tenets of evolutionary epistemology by focusing on related naturalistic approaches. In this section I shall directly start the analysis of evolutionary epistemology.

First of all, we need to point out that evolutionary epistemology is a variegated field in which different research programmes cohabit. Following Bradie (1986), it is useful to classify evolutionary epistemologies according to their programmatic aims by dividing them in two broad distinctive categories. EEM (evolution of cognitive *m*echanisms) is the evolutionary programme attempting to explain the development of cognitive mechanisms of living species, while EET (evolution of *theories*) is the programme whose agenda consists in analysing the growth of knowledge using an evolutionary model. These approaches are consistent. In fact, many epistemologists have endorsed both of them, Campbell (cf. sections 3.3 and 3.4) and Popper (1972) among them. Others have either embraced EEM (e.g. Ruse, chapter 2), or EET (e.g. Hull, chapter 4). EEM and EET share an adaptationist stance, but they differ on the postulation of the selective process of science. In this thesis I shall be critical towards EEM approaches and more sympathetic to EET ones. This is because my interest is mainly in developing a theory of scientific change, and, as I see it, EET is a much more promising approach to fulfil this aim than EEM. Of course opinions diverge.[36] More particularly, since the EEM programme accounts for the existence of knowledge by laying emphasis on the cognitive capacities of humans and their biological substrata (nervous system, brain, cortical cortex etc), it is largely irrelevant to the development of a theory of scientific change. As a consequence, I believe that what EEM has to say about the nature of science is either largely uncontroversial or false. After all, accepting the basic truth of Darwinism means endorsing the view that our cognitive abilities have been somehow shaped by the action of natural selection.[37] All evolutionary epistemologists must agree that the Darwinian approach to epistemological issues starts from the realisation that human beings are products of evolution and that their capacities for knowing have been honed by evolution (Bradie 1989 pp. 393-4).[38] What remains controversial is how much of human cognition can be explained in adaptive terms, how much is cultural, and how, supposing the evolutionary picture is true, such scientific information provides a justification of our knowledge. EEM does not try to fully answer such questions (which constitute the agenda of this thesis), while EET does.

[36] Radical EEM theorists (i.e. Ruse) contend that EET is not a satisfactory programme because there is no selective process operating in science that can be said to be entirely "analogous" to natural selection. I partly agree with this contention, but this does not mean that EET is not a viable research programme (cf. esp. section 4.7 and chapter 5).

[37] The truth of this thesis (a limited form of adaptationism) is not an issue at stake. In fact, there is a general agreement between EEM and EET theorists (and all sensible naturalists) on this score, as there is a general agreement between evolutionary biologists. This also explains why Campbell and Popper believed that a version of EEM must be part of EET.

[38] As a consequence, knowing, as a natural activity, should be treated and studied along lines compatible with its status, that is, via scientific method. Science is relevant to epistemology and epistemology is continuous with

It is difficult to identify features in common to all evolutionary epistemologies apart from general ones (e.g. the endorsement of the continuity thesis in some form, a Darwinian outlook, the commitment to an adaptationist stance). In fact, opinions between EEM and EET theorists differ drastically as soon as we try to detail the nature of the commitments. Evolutionary models provide diachronic epistemologies, in line with the historical character of the science of reference. But as soon as we try to articulate the historical nature of the models, we see again a variety of incompatible approaches emerging. For instance, while EET models are generally diachronic models of the growth of scientific knowledge, EEM models are directly concerned with the "evolution" of the cognitive structures of the human species and only indirectly (if a link can be established at all) with the "evolution" of science. While EEM's science of reference is evolutionary psychology, in EET's case the history and sociology of science play a fundamental role.

Perhaps we might be able to identify common traits among evolutionary models if we consider the relationship between them and traditional epistemologies. Campbell (1977) argued that the relationship could be defined in terms of competition, succession and complementarity. In the first case, evolutionary models are trying to address the same problems as traditional epistemologies and offering competing solutions to them; in the second case evolutionary models do not address the same questions, because the problems of traditional epistemology are either irrelevant or unanswerable; while in the third case evolutionary and traditional models are to be seen as having different agendas and as complementary disciplines. It is difficult to pinpoint a common thread here too. As will become familiar in the following pages, the evolutionary models I shall illustrate in the thesis do not propose to abandon the traditional normative agenda, by completely substituting the normative with the descriptive. However, in general EEM models are more tradition-bounded, while EET ones try to revise the traditional normative project.

science. In other words, the Darwinian tenet is central to justify the continuity thesis at the heart of the entire naturalistic approach to epistemology.

Even though I tend to think that we should consider evolutionary epistemologies as pursuing projects that are complementary to those of traditional approaches, in other ways it is difficult not to see certain evolutionary models (like other naturalistic, but not evolutionary, models) as successors to traditional approaches. In fact, evolutionary models inevitably beg some traditional concerns. In any case, the problem of validation remains central for all the evolutionary models I shall illustrate. This is the reason why I think it is better to see evolutionary epistemology as complementary to traditional approaches.[39]

Evolutionary epistemologies face the same basic challenges as other naturalistic approaches. One of the aims of my thesis is to assess how the evolutionary approach to epistemology contributes to answering the challenges typical of epistemic naturalism illustrated in section 1.2 (cf. chapter 5). In the rest of this section I give an outline of this contribution.

Starting with circularity, the most relevant contribution to the dissolution of the circle problem coming from evolutionary epistemology pertains to the endorsement of Darwinism. The interesting issue regards the nature of the Darwinian commitment. Darwinism means at least two different things to different evolutionary epistemologists. First, there are those who understand Darwinism literally and who interpret the significance of evolutionary theory in a purely biological fashion. Second, there are those who consider Darwinism as more than a purely biological theory and who consider evolutionary theory as a general theory of cognition. Typically, EEM theorists endorse the first vision, while EET theorists' commitment is more substantial (compare chapter 2 with 3). One common feature of the two Darwinian commitments is that they are supposed to function to stop the circle argument, even though this effect is achieved in different ways. In the first case, Darwinism stops the circle by assuming that the epistemic role and value of evolutionary theory cannot be disputed and put into question. In this way, it provides the "empirical foundation" to rebut the sceptic, via the argument from natural selection, which maintains that humans' cognitive endowment

[39] Thus, to claim that evolutionary epistemologies are alternatives to, say, synchronic (e.g. formal) models of rationality is doubly misleading. First, because, for instance, some formal epistemologies (e.g. Bayesianism) are not

must be minimally rational otherwise our species would not have survived so long (cf. section 2.4). In the second case, Darwinism stops the circle by postulating that the only way to acquire knowledge is via the basic Darwinian process of blind variation and selective retention.[40] This posit can either be seen as an empirical hypothesis or as a methodological tenet. The first move is in line with naturalism but a difficult hypothesis to establish.[41] The second move is clearly compatible with the naturalistic commitments of evolutionary epistemology with perhaps the advantage of capturing the best of both the "dogmatic" (a.k.a. "not under discussion at the moment") and hypothetical character of the basic selectionist hypothesis (cf. section 5.2.2).

On the issues of normativity and relativism the positive contribution of the evolutionary approach is more significant and innovative.
The issue of normativity remains central for any evolutionary epistemology. Like all naturalistic models, evolutionary epistemology consists of a descriptive and of a normative part. Evolutionary epistemology pursues the normative agenda by basing its normative claims on the descriptive ones. Since evolutionary epistemologies are committed to the pursuit of the meliorative project of improving the effectiveness of the scientific process, the norms extrapolated are justified only insofar as they are epistemically conducive. Of course, evolutionary models of the evolution of science can be criticised for reasons that are parallel to those adduced against many other naturalistic models. As typically pointed out by traditional epistemologists, one of the limits in the attempt of deriving norms from facts is that these facts are already "coloured" by prior philosophical commitments. To this contention the evolutionary epistemologist, like any fellow naturalist, should reply that the philosophical

incompatible with historical approaches (as already noted by Salmon - 1990). Second, because formal and evolutionary epistemologies might, as a matter of fact, be complementary approaches.
[40] The best formulation of this argument is found in Campbell (cf. sections 3.3.and 3.4). To illustrate the nature of this central tenet of EET (with which the reader will become soon familiar) it will be sufficient to recall Popper's famous position according to which trial and error (i.e. the process of the emergence of biological or cultural variants and their subsequent selection) is the epitome of the genuine scientific approach (i.e. it identifies the logic of science).
[41] As Cziko (2001) clarifies, the universalism of this Darwinian predicament at the basis of universal selection theory has to be circumscribed in order to make it significant (cf. also 3.3).

theory on which they rely to pursue their normative endeavour is equivalent to the theoretical part of any science, part and parcel of the empirical theory of science they try to construct.[42] Focusing on EET models, my aim is to develop EET in two normative directions. First, as will become clear after reading chapters 3 and 4, EET can be seen as a social epistemology analysing science as an adaptive social system governed by peculiar institutions and social norms that are designed in order to govern the achievement of its institutional goals. EET bases its analysis on descriptive knowledge of many kinds (e.g. biological, historical, sociological) and in return allows epistemologists and science-policy makers to extrapolate and devise recommendations in order to render science a more effective epistemic practice. Second, EET attempts to provide a general analysis of selection processes by identifying the conditions that render selection processes adaptive (i.e. aim conducive). Science as a selection process can be thus analysed as an adaptive system from this more general, and fully evolutionary, perspective. These issues will be treated in detail in section 5.1.

Evolutionary epistemology's contribution to the rebuttal of the relativistic challenge will be analysed in section 5.3. This contribution centres on the constructivist thesis embraced by evolutionary epistemology. Let us see how this thesis is articulated.

Consider that the central epistemological problem of evolutionary epistemology is to study how knowing organisms interact with their environments in order to produce knowledge. It goes without saying that the picture of reality an organism constructs depends on what kinds of raw materials it has available (e.g. cognitive structures, observational powers, scientific notions etc.). Thus, first of all, and rather trivially, the epistemological relativism of evolutionary epistemology is due to the intrinsic impossibility to directly know reality: knowledge is always an edited, indirect and fallible version of reality as it always amounts to a compromise between the characteristics of the vehicle of knowledge (e.g. a particular organism or group of organisms) and referent properties. Each kind of knowing organism constructs a picture of reality that is partly and ineluctably based on its needs and capacities (I call this

[42] In this sense I completely agree with Rosenberg's position (cf. note 20).

minimally relativistic position “perspectival relativism”). But such epistemological relativity (be it species-relative or socially constructed depending on what kind of vehicle is taken into account) is not per se evidence to condemn evolutionary epistemology as dangerously relativistic.

But there is an additional and more significant sense in which evolutionary epistemology, by embracing the naturalistic approach, can be seen as relativistic. This sense pertains to the hypothetical, contingent and local nature of the normative claims generated by evolutionary epistemology. Once the psychological, biological, sociological, historical and generally cultural causes of the origin of belief, and more importantly its function, have been studied, there is no further question to be asked about validity. Relativism is a natural outcome of the instrumentalism and lack of universalism of the naturalistic turn.

So, how does evolutionary epistemology defeat relativism and save the objectivity of science? The contribution of evolutionary epistemology[43] to rebut relativism concerns the articulation of the notion of epistemic norm or standard. Epistemic standards can be seen, from the evolutionary perspective, as evolutionary products, where by evolutionary product I mean product of selection processes of some kind. In science such processes are not biological, but cultural and social. Epistemic standards can be more aptly seen as social products and can be studied by focusing on the details of their cultural emergence and on the function they perform in the scientific game. The advantage to see norms as evolutionary products pertains to them being treated as historical objects that do not belong to a particular perspective. Rather, the epistemic standards studied by evolutionary epistemology have an autonomous life and a community-transcendent nature. In this way the objectivity of science is not dependent on the judgement, for instance, of a particular community of trained scientists, but it is constrained by the causal role that some multi-generational standards play in scientific practice. As Campbell put it (1981), the objectivity of science depends on the disputations between multigenerational “communities of truth-seekers”.

[43] I here only refer to the EET models I favour.

In the next chapter I shall focus on the EEM model of evolutionary epistemology developed by M. Ruse, while from chapter 3 I will focus on EET models. The next chapter will be mainly critical. The conclusive chapter, which deals in detail with the issues treated in this section, will be based on the analysis of the evolutionary models illustrated in chapter 3 and 4, that is, on the most original, interesting and articulated evolutionary models of scientific change in my opinion ever proposed (respectively by Donald T. Campbell and David Hull).

Chapter 2: The Case Against EEM

"It is possible to imagine that chimpanzees have an innate fear of snakes because those who lacked this genetically determined property did not survive to reproduce, but one can hardly argue that humans have the capacity to discover quantum mechanics for similar reasons."
Noam Chomsky, "The Managua Lectures", p. 156

In this chapter I shall assess and criticise the EEM programme in evolutionary epistemology, which accounts for the existence of knowledge by putting emphasis on the cognitive capacities of humans and their biological substrata (nervous system, brain, cortical cortex etc). I believe that EEM is a largely unsatisfactory epistemological approach, which fails to describe and explain adequately the nature of the scientific selective process. In this chapter I shall only indirectly show, by individuating EEM's deficiencies, that EET is a much more interesting and epistemologically relevant research programme than EEM.

I shall now present Michael Ruse's EEM approach and critically analyse it. I chose this model because I believe that it epitomises vices and virtues of any EEM approach. From the assessment of this epistemology I will conclude that any EEM programme constructed along the same lines lacks the sufficient epistemic resources in order to provide a realistic account of the nature of science.

2.1 Ruse and the sociobiology of science

In the book "Taking Darwin Seriously" Ruse claims that, if we are to take Darwin seriously, if we are to give a Darwinian account of the nature of science, we must ultimately refer to the cognitive units that regulate the working of the human mind. These cognitive units are the "epigenetic rules", developmental regularities whose existence has been posited by two eminent sociobiologists (Wilson & Lumsden 1981). Ruse argues that the epigenetic rules are of primary epistemological importance because some of them are the psychological instantiation of rational methodological rules.[44]

Ruse's account of evolutionary epistemology is based on the belief, according to him universally shared, that there is a specific methodology that "creates" science, and that the epigenetic rules lie behind this methodology.[45] The doctrine is eminently sociobiological: since there are rules setting up incest barriers, why should we deny there are rules for approval of modus ponens? This is the sociobiological hypothesis. I shall criticise Ruse mainly because I think that the epigenetic rules Ruse posits do not provide the foundation of science. As a result, Ruse's EEM is epistemologically irrelevant in many respects (cf. section 2.5).

In order to do so, I need to illustrate the nature of the sociobiological hypothesis. This is an empirical hypothesis concerning what can be termed the "science-forming capacity", that is, that bunch of cognitive processes that are necessary and sufficient in order to have science and to fully explain its nature. Some epigenetic rules constitute such science-forming capacity. These rules are innate or genetically encoded (cf. section 2.2), and they are adaptations (cf. section 2.3). Furthermore, such rules are reliable and somehow justifiable (cf. section 2.4). What this hypothesis means is that science can be completely characterised in terms of these rules. As a consequence, Ruse's EEM is a complete account of what is epistemologically relevant.

In order to assess this hypothesis, we need first to assess what Ruse says about the nature of epigenetic rules. The epigenetic rules that, according to Ruse, constitute the science-forming capacity are the laws of logic (in particular modus ponens), a rule for induction and a rule for

[44] This is an unchallenged tenet of Ruse's EEM, since he assumes that there is a somehow fixed methodological recipe to achieve knowledge. Some of the problems of this view could be thus highlighted. First, an analysis of the practice of science shows that even "good" scientists fail to instantiate the optimal mental processes corresponding to the optimal methodological rules of scientific method Ruse posits. Second, assuming that optimal rules exist is in itself a doubtful claim. Third, even though such rules exist, it is in practice very difficult to articulate them properly and fully, given their context-dependence, environmental-relative and "local" nature. Fourth, even if there were a fixed methodology, it would be insufficient to dictate what we "should" believe, because what to believe is always underdetermined by good methodology.

simplicity. In order to support the methodological significance of modus ponens as an epigenetic rule Ruse presents the following argument:

"Consider two would-be human ancestors, one with elementary logical and mathematical skills, and the other without very much in that direction. One can think of countless situations, many of which must have happened in real life, where the former proto-human would have a great selective advantage over the other. A tiger is seen entering a cave that you and your family usually use for sleeping. No one has seen the tiger emerge. Should you seek accommodation for this night at least? How else does one achieve a happy end to this story, other than by an application of those laws of logic that we try to uncover for our students in elementary logic classes."[46]

Ruse goes on to provide similar kinds of arguments in order to establish the adaptiveness and *prima facie* rationality of all the basic rules constituting his basic methodology. For induction he quotes Quine:

"Creatures inveterately wrong in their inductions have a pathetic but praise-worthy tendency to die before reproducing their kind."[47]

He then goes on to justify simplicity:

"....we can see good biological reasons for favouring simplicity....the primate who innately favoured the simpler option [between two evidentially equivalent ones] would go ahead of his/her fellow with a taste for the complex. The former would be spending a lot less time and effort on his/her decision-making and execution."[48]

Finally he shows that any philosophical and critical attitude cannot be selected for by evolution:

"One hominid arrives at the water-hole, finding tiger-like footprints at the edge, blood stains on the ground...[etc]. She reasons: 'Tigers! Beware!' And she flees. The second hominid arrives at the water, notices all

---

[45] M. Ruse (1986) pp.160-169.
[46] ibid. p. 162
[47] ibid. p. 162
[48] ibid p. 162

of the signs, but concludes that since all the evidence is circumstantial nothing can be proved. 'Tigers are just a theory, not a fact'. He settles down for a good long drink. Which of these two hominids was your ancestor?"[49]

These arguments should provide a positive case "favouring the claim that scientific methodology is grounded in epigenetic rules, brought into existence by natural selection."

The epistemic status of the sociobiological hypothesis must be assessed. In this respect, Ruse claims that his sociobiology of science is an empirical hypothesis already empirically supported:

"...there is already empirical evidence....backing the epigenetic-rule status of the formal and informal aspects of scientific methodology."[50]

The evidence comes from many empirical disciplines. Comparative anthropological studies of different societies show that:

"...beneath differences [cultural and individual], the same kind of generalizing ability seems to be at work, with people of all kinds drawing similar-type connections...between diverse phenomena.....All humans share...recognizable (to us) standards, like the refusal to take coincidences as coincidences, preferring rather to look for underlying unifying causes."[51]

Ruse dismisses any kind of cultural relativism by basing his analysis on the assumption (or, better, working hypothesis) that "human nature" is an epistemologically relevant posit. The hypothesis of the psychological unity of mankind is at the heart of Ruse's EEM (cf. section 2.2). What Ruse also argues for is that the cognitive uniformity he posits needs an explanation in terms of "shared adaptive needs", that in the environment of ancestral adaptation such needs were pressing, given that "it is unlikely that nature would have left them open to random environmental disturbance". Epigenetic rules are therefore seen as adaptations. Natural

---

[49] ibid. p. 163
[50] ibid. p. 164
[51] ibid p. 164

selection should start to make innate "the inclinations which would prove of maximum selective value." However, the case for innateness can neither be solely based on the link between advantageousness, nor only on anthropological evidence. Therefore, the innateness thesis is based on further empirical evidence coming from developmental psychology. Language acquisition studies show that children have innate dispositions to think along certain ways, that they do not "learn" how to be language users:

> "A second piece of evidence pointing at the reality of the supposed epigenetic rules of science comes from studies of childhood development. As with language, it is becoming more and more evident that children do not simply learn the crucial elements of their cultures, as though their brains were *tabulae rasae*, soaking up any piece of information offered, in the order it is offered. Learning occurs in highly stylized ways..."[52]

Ruse points to further evidence to support his view. Many cross species studies in comparative zoology and ethology show that even primates have rudimentary notions of logic. All this evidence supports Ruse's suspicion that the rules of logic, induction and simplicity he believes constitute the method of science are epigenetic rules, since they are essential to "get through life". Such rules are "ingrained in our nature", that is, part of our "biologically informed" and "innate conceptual apparatus". I find Ruse's claims very suspicious for many reasons that I shall now illustrate.

2.2 Innateness and species-typicality

In section 2.4 I shall assess the normative value of Ruse's hypothesis. For the moment, however, I shall confine my treatment to the descriptive value of the sociobiological hypothesis, which is far from established. What does it mean, in particular, to believe that the epigenetic rules governing scientific behaviour are genetically based, innate, species-typical and adaptive? I contend that all these claims are dubious.

---

[52] ibid. p. 165

In order to assess the various tenets of Ruse's hypothesis, let us focus on just one among the many methodological rules Ruse believes provide the skeleton of science. What does it mean to consider, for instance, modus ponens as an epigenetic rule? If modus ponens is an epigenetic rule, it is an innate, species-typical, and adaptive trait. I shall start by assessing the claim that modus ponens is innate, and then move to the universalistic hypothesis.

In Ruse's view, modus ponens is an epigenetic rule, where such rules are developmental regularities. In this sense, the concept of innateness adopted by Ruse is that of developmental fixity: a trait is innate when it is developmentally canalized, that is, when its development is uninfluenced by environmental interferences, when the trait will emerge in all "normal" developmental environments. There is evidence to consider modus ponens developmentally fixed, but this does not mean that modus ponens is genetically encoded. In general, phenotypic invariance across "normal" environments or insensitivity to environmental factors could be either determined by the fact that the trait is genetically programmed, or by the fact that the causally relevant factors affecting development are invariant across normal environments. In this sense, modus ponens could be developmentally fixed because of the social and cultural features characterising normal environments. That is, the same cultural environment can induce the same developmental outcome without there being genetic encoding. Ruse claims that modus ponens is universal, but this could mean that the same cultural conditions determine the same developmentally fixed outcome. So cultural evidence of this kind cannot be used to show that modus ponens is genetically encoded. The developmental evidence Ruse presents does not adjudicate in this case because modus ponens could develop in all "normal" or "healthy" children because of social factors. For instance, Gopnik argues that children do not seem to be able to use modus ponens until they develop a so-called theory of mind (Gopnik 1996). Autistic children who fail the false belief test might not use modus ponens, since this test requires using such cognitive strategy. But this could be due to the fact that the social environment in which autistic children grow up is not socially rich enough. Ruse also claims that the thesis of phenotypic universality can be extended to other species, as evidenced by studies of comparative ethology. But even in this case the cultural hypothesis is not disconfirmed, given that the chimps to which Ruse refers live in social environments. The evidence showing that modus ponens is developmentally fixed is compatible with the hypothesis that modus ponens is genetically encoded, but it is not sufficient.

What we need to assess is whether modus ponens is genetically encoded. Is there any evidence in this sense? Ruse claims that evidence coming from cultural studies shows that modus ponens seems to be a phenotypic universal. One problem with this view is that modus ponens could be universal without any genetic encoding, as are some traits that are universal across cultures (e.g. belief in the existence of some transcendental entity). For this reason, inferences from universality to genetic encoding are suspicious (Griffiths 2002). Further evidence concerns the existence of genetic conditions that impair the use of modus ponens. Perhaps autistic children, unable as they seem to be to develop a theory of mind, might be impaired in this sense. The issue remains whether autism is a genetic deficiency. The evidence in this respect is not, however, conclusive. The best argument Ruse presents concerns the advantageousness of modus ponens. In this sense, Ruse could infer innateness from adaptiveness: given the selective advantage determined by modus ponens, then "if selection is going to make anything innate, it would do well to start here." In order to test this adaptive hypothesis (cf. section 2.3) we would need to show that the phenotypic variation in the ancestral population was under genetic control. Only if this is the case natural selection can encode modus ponens as an obligate response. In general, whether a cognitive trait is encoded as a cognitive mechanism depends on the benefits in terms of reliability, informational economy and speed of processing it determines. As far as informational economy is concerned, having modus ponens encoded is not necessarily fitter than learning to use it. This is because, in general, encoding a cognitive strategy is not necessarily more adaptive than learning, given that improvements in fitness can come both by adding and reducing cognitive equipment (Sober 1994a). At this point, one issue is whether it is better to encode modus ponens as a facultative (e.g. use modus ponens when you see predators), or as an obligate response (e.g. use modus ponens in all domains). Another issue concerns reliability. The point I want to make is that, in any case, learning can be more economical than encoding. If the learning process is social, then teaching how to use modus ponens is more economical than encoding it as an obligate or facultative response, given that genomic space is saved. Thus, learning how to use modus ponens might be more adaptive (i.e. informationally fitter) than having it encoded. This points to a general limitation of Ruse's model of evolution, where only genetic encoding makes the trait heritable.[53] In this respect it might be noted that most cognitive variation encountered in humans is not genetically encoded. For instance, language diversity is most likely due to environmental factors, not to genetic ones. In the case of modus ponens, we need to assess whether the ancestral variation was determined by environmental rather than genetic factors. In the former case modus ponens could be inherited culturally.

But questions of informational economy are just one factor to be considered, given that reliability considerations have as much importance. The question is whether having the trait learned or innate is more reliable. Given the reliability of modus ponens, we can accept Ruse's hypothesis. This is because modus ponens is the most reliable strategy across all contexts and potential environments. It is generally argued that in stable environments it might be better to have behavioural responses as hard-wired as this will save on the costs and risks of learning behaviour later in development (Williams 1982). So the question we need to ask is whether the environment of ancestral adaptation was stable as far as the trait modus ponens is concerned. It seems to me that it is difficult to envisage an unstable environment as far as modus ponens is concerned. This is because, again, modus ponens seems advantageous in all environments. Thus, even if questions of informational economy might not adjudicate the issue, it seems to me correct to hypothesise that having modus ponens as encoded as a cognitive strategy seems to give benefits, since such encoding determines no loss in mental plasticity and general adaptability to different environmental conditions. So Ruse might be right in inferring encoding from adaptiveness as far as modus ponens is concerned.

[53] Even though, in principle, Ruse might accept the cultural transmission of biological adaptations, according to an

It is now time to extend the innateness argument to the other cognitive strategies constituting the science-forming capacity. The assessment of the innateness hypothesis is even more difficult in this case, basically because the only good argument Ruse proposes (based on reliability considerations) cannot be extended to inductive and simplicity rules. In fact, inductive and simplicity rules are not as reliable as modus ponens. Their reliability is constrained by questions concerning the domain of cognitive application of the rules, and the biological advantage is much less clear. This criticism becomes stronger the more we ascend the level of complexity of the cognitive strategy, given that, for example, we cannot assume that the environment of ancestral adaptation was stable as far as the highly context-dependent cognitive strategies peculiar of science are concerned (Dupre' 2001, chapter 2). What Ruse proposes is, however, an argument about reliability that has limited scope. Modus ponens is peculiar in this sense. But an argument is not "evidence" in favour of the innateness hypothesis. Ruse should seriously test innateness against the hypothesis about cultural inheritance, but he does not. As a consequence, we have no reason to believe that even modus ponens is innate.

But even if we grant that modus ponens is innate, and that all other cognitive strategies constituting the science-forming capacity are, this does not mean that, as Ruse claims, we are somehow "born" with the capacity of doing science. The obvious reason for this is that the cognitive strategies identified by Ruse are insufficient to produce science as we know it. A further reason is that to have science we do not need science-typical cognitive invariances across people. That scientists behave in such multifarious cognitive ways is generally considered a maladaptive feature of the scientific process (at least as far as scientific selection is concerned). Such maladaptive behaviour can be both explained by pointing out that the cognitive strategies used by scientists were adapted to the Pleistocene, and that they are for this reason not good enough for the scientific environment. But it can also be due to the fact that a diversity of human psychologies has been maintained by natural selection (Sterelny

---

extended model of evolution by natural selection (cf. section 3.2)

1995). Therefore scientists' behaviour is idiosyncratic in scientific matters because humans are not designed in the same way. One reason for this is that it is difficult to argue that there must have been a single best cognitive design to solve the cognitive problems faced by hunter gatherers. In the scientific case matters are even worse, since the notion of single best cognitive design necessary to solve the vagaries of scientific problems faced by scientists does not make any sense. In brief, human cognitive differences might not simply be the result of developmental vagaries, but rather evidence of the fact that there is no psychological unity of mankind after all. As a result, science's success might be explained by showing that cognitive variation is an essential element, rather than a negative by-effect, of the scientific process (both as far as scientific discovery and selection are concerned - cf. sections 4.4 and 5.2.3).

The thesis of the psychological unity of mankind states that the human species is cognitively uniform, and that there exist some universal and species-typical traits. As Ruse puts it:

"For the Darwinian, the required universality follows on the unity of humankind."[54]

Some "Darwinian" philosophers have expressed doubts about the notion of human nature. They have both expressed doubts about its ontological status (the concept of human nature belongs to the vocabulary of folk essentialism and has no reference – Griffiths 2002), and its explanatory role (since Darwinism is a populational theory the notion of human nature or "genotypic/phenotypic normality" has no value to explain our, for instance, moral behaviour – Hull 1986). Others have presented empirical evidence that challenges the assumption that there are psychological invariances across people (Stich 1985). Such evidence shows that there are significant differences in cognitive performance across individuals, and that the amount of cognitive variation among humans should not be underestimated. The quick and dismissive reply to this contention usually takes two forms. The first is to explain away human cognitive idiosyncrasies by arguing that the evolved universal design developed in different ways

[54] M. Ruse (1986) p.189

because of variable environmental factors (the strategy adapted by sociobiologists and evolutionary psychologists). The second is to posit a distinction between performance and competence, that is, starting from the supposition that human beings are endowed with an almost perfect "rational" cognitive capacity and that their mistakes are just caused by lack of attention and external interferences (Sober 1981). Both arguments might be used to defend the thesis that the science-forming capacity is unique, but cannot be used as a priori recipes to save the universalistic assumption. For instance, Ruse could point out that the science-forming capacity is an informationally encapsulated module like the language organ. However, in the case of language we know that severely cognitively-impaired children are good language users nonetheless. But in the case of the science-forming capacity, we simply do not know whether children impaired in other cognitive skills might become extraordinary scientists. For this reason, there is no evidence to believe that the science-forming capacity is a module.

I conclude that there is no evidence that can adjudicate between Ruse's hypothesis about our inborn ability to engage in science, and the alternative hypothesis that such capacity is not innate at all.

2.3 The adaptive hypothesis

In this section I will not criticise the adaptationist stance of Ruse's EEM in general. I am an adaptationist as well, as any evolutionary epistemologist is. As a matter of fact, I think adaptationism is a much more general hypothesis than Ruse believes (cf. sections 3.2. and 3.3). In this sense, I believe that EEM and EET do not endorse a "Panglossian" or naive adaptationism (Gould & Lewontin 1979). Adaptationism is compatible with the view that other causal factors apart from selection (e.g. drift, linkage of traits, accidents etc) influence evolution. Adaptationism does not assume a priori that all traits are the result of a selection regime, and that there are no "spandrels". Adaptationism accepts that some constraints limit the optimising power of selection processes. Adaptationism does not provide "just so stories" or untestable hypotheses. The consensus is that adaptationism, when purged of its naivety, is the fundamental working hypothesis in evolutionary theory. Having said this, we need to see whether Ruse establishes that the science-forming capacity consists of adaptations. Again, let us start by considering whether modus ponens is an adaptation, and then move on to the other epigenetic rules. In this section I shall assume that there is some genetic encoding for modus ponens. The question to be assessed is whether the trait is an adaptation.

The first point to stress is that inferences from innateness to adaptiveness are not always correct. This is because there exist traits that are innate (i.e. developmentally fixed) but not adaptations (Griffiths 2002). Modus ponens could be just one of these traits. For a trait to evolve by natural selection we need heritable phenotypic variation in fitness within the ancestral population.

The variation condition seems satisfied. There is definitely between species variation. Ruse claims that chimps, our nearest relatives, use modus ponens; however, not all mammals do. The adaptive hypothesis could thus be tested by taking into consideration such phylogenetic information. I believe there might be also within species variation. Even though modus ponens is a universal trait, perhaps in some contexts people find it difficult to use it, or perhaps some cognitively-impaired people cannot use it (e.g. people who fail the false belief test, which presupposes the use of modus ponens). The point to be stressed at this juncture is that it is not necessary to have phenotypic variation in the actual population, since the trait might have gone to fixation.

The second condition concerns heritability, which is by assumption satisfied, since in this section I grant that modus ponens is genetically encoded (despite the critical comments made in section 2.2).

We need to assess if the third condition is satisfied, that is, whether modus ponens is a trait that gives a selective advantage in biological fitness. In this respect, Ruse seems to presuppose that modus ponens has evolved because it conferred a biological advantage. However, it might be argued that modus ponens is rather the adaptive outcome of a process of cultural selection of some kind. For instance, it might be argued that humans use modus ponens because our ancestors knew that using this cognitive strategy was culturally advantageous (e.g. in the sense that it helped them in coping better with the environmental problems they faced). This cultural hypothesis could be also coupled with a group selection one: primate groups that used modus ponens were culturally fitter than those that did not use it. As a consequence, modus ponens became a developmentally fixed trait because modern humans live in human groups that have inherited, culturally, such a penchant for modus ponens. Such hypotheses could be further confirmed if it is shown that chimps are capable of "culture".

Ruse is correct in stressing that modus ponens is biologically advantageous. We could hypothesise that one mutant primate eventually reproduced more because of the selective advantage given by the novel trait modus ponens. In order to test such an adaptive hypothesis we need to reconstruct the nature of the ancestral environment. In this sense, we would, for example, need information about the cognitive structure of our ancestors, and about the cognitive tasks that they performed relative to their environmental needs, and then see whether this information provides reasons to believe that modus ponens is an adaptation. I agree with Ruse that modus ponens is such a reliable trait that natural selection will, assuming that phenotypic variation is under genetic control, certainly try to make it innate. Modus ponens is advantageous in all environments, plus it is difficult to imagine a cognitive alternative of equal fitness. Therefore, modus ponens seems to me a good rule both for avoiding tigers (cf. Ruse's argument in section 2.1), and for scientific purposes. Ruse might be correct in thinking that modus ponens is an adaptation, but the problem with this view is that there are many biologically advantageous traits that are not adaptations. Humans are unable to synthesise vitamin C. It would be advantageous to have such an ability, but "genetics can get in the way" (Sober 2000 p.125). Even if making a fire is a biologically advantageous trait, it is not an adaptation because some vicarious selective process has substituted the action of natural selection guaranteeing some kind of "inheritance" (cf. section 3.4). Given social learning there is no need for natural selection to encode the trait.

Another element to consider when testing the hypothesis is the nature of the constraints limiting the optimizing action of natural selection. One of these is the amount of putative genetically encoded cognitive strategies among which modus ponens was chosen by selection. We need information about the amount of genetic variation encoding for different cognitive phenotypes present in the ancestral population. Natural selection can only choose between available traits. In this sense, the number of alternatives upon which natural selection operates might be vast, even though it is certainly not infinite. After all, the fact that zebras have not developed guns to kill predators is just evidence of the fact that, even in the long run, evolution is not miraculous. We have no information of this kind.

Another constraint is given by the already operative effect of more general or complex cognitive strategies. In this case modus ponens could be a spandrel or by-product of the selection of some other cognitive strategy, for example the ability to use language. However, given that even chimps seem to be able to use modus ponens, we could hypothesise that it did not evolve as a by-product or spandrel of the selection of the "language organ". But the question remains open as to whether it is a by-product of the selection of some other trait. There is no evidence to test (comparatively) Ruse's hypothesis in this respect.

My point so far has been that testing an adaptive hypothesis is extremely intricate, especially if we lack the relevant information. Ruse might be correct in thinking that modus ponens is an adaptation, but the evidence he puts forward is insufficient to adjudicate between his hypothesis and the alternatives.

I now want to move to another issue. Even if we grant that modus ponens is an adaptation, what does this mean as far as the other cognitive strategies constituting the science-forming capacity are concerned? First, if modus ponens is an adaptation this does not imply that induction is. With this I am not contending that, as Russell's chicken demonstrates, induction is not a biologically advantageous trait (even though one might doubt its reliability). In my opinion, the adaptive hypothesis becomes even more dubious the more we ascend the level of complexity of the cognitive strategy used by scientists.

Second, even though modus ponens, induction and simplicity are adaptations, it might be wondered whether such a science-forming capacity can get science off the ground. I am not disputing that Ruse's epigenetic rules are necessary for science, but simply that they are not even remotely sufficient neither to explain science's success, nor to describe scientists' cognitive behaviour. This is because a cognitive strategy good for hunter-gathering is not per se a good cognitive strategy for science. This could be so only if we show that the environment of ancestral adaptation poses the same kind of environmental problems that we find in science.

Ruse's science-forming capacity identifies the essential rudiments that make science possible. In order to criticise Ruse's EEM it could be sufficient to ask why, if both humans and chimps can use modus ponens, chimps have no science.

2.4 The argument from natural selection

In sections 2.2 and 2.3 I showed that Ruse's sociobiological hypothesis is far from established as far as its descriptive value is concerned. Furthermore, even though modus ponens happens to be innate, genetically encoded, species-typical and adaptive, it does not follow that all the other traits constituting the science-forming capacity are. But let us assume that Ruse's descriptive hypothesis is correct, and that the science-forming capacity exists. Let us also assume that reference to such capacity explains all there is to explain about science. What we need to consider now is whether we can equate biological processes with epistemic and rational knowledge-gaining. What I want to assess is whether Ruse's EEM has any normative value: how can the epigenetic rules Ruse posits be justified and validated? For this purpose, Ruse relies on the argument from natural selection, which can be presented, schematically, as follows:

a) Hypothesis about biological evolution

Epigenetic rules are the products of biological evolution.

b) Hypothesis about the units of selection

Evolution is caused by natural selection acting at the level of individuals (and genes). This means that the epigenetic rules are individual-level adaptations. Epigenetic rules are psychological instantiations of the epistemic standards governing the scientific process.

c) Optimisation hypothesis

Natural selection will eventually select for the most fitness-enhancing cognitive strategy in the gene-pool, by choosing between a massive number of alternatives, therefore picking an approximation to the optimal cognitive strategy in the long run.

d) Hypothesis about reliability

The selected and fittest cognitive strategies that approximate the theoretical optimum are truth-conducive. Well-designed cognitive systems are reliable because reliability is favoured by natural selection, since true beliefs are more adaptive than false ones, in the sense of enabling the organism to cope better with the necessities of the environment.

From premises (1) to (4), this conclusion follows:

(C) Human cognitive strategies, by being reliable, are rational and justified, that is, well-designed organisms with well-designed cognitive systems have rational cognitive systems.[55]

Even though I am quite sympathetic to arguments of this kind, but of different form, I believe that this particular version does not succeed in establishing the truth of the conclusion. My stance seems contradictory. I believe that the argument from natural selection is a good attempt to justify standards. This is because I believe that naturalistic attempts are the only available ones, given that the a priori path is forbidden to evolutionary epistemologists (cf. section 1.1 and 1.2.1). I believe that naturalistic arguments of this kind might be used for normative purposes, and that, naturalistic fallacy notwithstanding, their normative worth should be assessed carefully. What is wrong about this argument is that the truth of the

[55] An illustration of the argument can be found in Dennett (1982).

premises is far from established. This is because the argument is based on premises and assumptions that are not substantiated, and whose descriptive accuracy is hardly undisputed.

Premise (a) is based on the hypothesis that biology is more important than culture. However, biological factors are not the only ones responsible for shaping our cognitive strategies. Cultural evolution has certainly aided biological evolution in shaping human cognitive design. This fact of the matter is a well-established truth that cannot be denied by sociobiologists, since culture is under the influence of a mechanism of selection that causes the spreading of behaviours that do not maximise biological (e.g. genetic) fitness (cf. section 3.2).

Premise (b) might also be mistaken, especially if the alternative adaptive hypothesis of group selection is not tested properly, but simply rejected a priori. In this sense we need to consider that the mind might be at least partly a group-level adaptation instead of an individual level one, especially if it is shown that it helped groups to respond better to environmental challenges, for example because they facilitated cooperative and altruistic behaviour.

As premise (c) is concerned, I believe that even if we grant that the science-forming capacity is an adaptation, it does not follow that evolution by natural selection constructs organisms that are optimally well designed to the environment, that in the long run, given a certain constancy in the environment, natural selection will come to choose an organism that is a close approximation to the cognitive theoretical optimum. This is an extreme form of naive adaptationism that cannot be endorsed unless we have a very unrealistic view of the optimising capacities of natural selection.

The assumption that natural selection optimises must be relativised to the nature of architectural and developmental constraints (cf. section 2.3), on the amount of genetic variation present, on the time allowed to natural selection to operate, and on the nature of the environment. The problem with genotypic availability is that the most reliable and rational cognitive strategies may well not be among those available to natural selection to select for. To assume the contrary is to assume that all the possible cognitive strategies have been genetically instantiated at some point or another during the course of evolution, which is absurd. A more optimistic view concerning the amount of genetic variation is taken by Dawkins (1999 pp. 43-50), who claims that the amount of genetic variation must generally be very large given what we are able to do by means of artificial selection. This argument is criticisable given that artificial selection is not always successful (e.g. we cannot breed cattle with a bias towards producing heifer rather than bull calves). But Dawkins' point is hardly applicable a priori to the cognitive realm. The second constraint concerns the nature of the time lag allowed to selection to optimise. It might be argued that, given sufficient time, natural selection can overcome almost any constraint. For example, if a gene has one beneficial and one harmful effect, there is no reason not to believe that a mutation could arise to detach the two phenotypic effects (Dawkins 1999 p. 35). Again, the issue is empirical and it cannot be said that, in general, any cognitive trait that is conducive to survival will, over a long period of time, come to fixation.

An argument against the naive optimality assumption is that natural selection favours "local" optima, because, being a blind mechanism, it cannot determine whether there are "global" optima. Natural selection cannot select for global optima because that would require knowledge that selection has not and cannot have, not being a cogniser. It would require information regarding environmental change, a complete forecast of the future that is unavailable except to an omniscient God. In addition, it is not so easy to argue that there exist global optima (apart from modus ponens). Furthermore, the local optima are possibly many. This points to another argument against Panglossian adaptationism, namely, the existence of multiple adaptive peaks. The point is basically that, relative to a particular environment, there are many optimal relationships of fitness enhancement. To illustrate the case consider the Wimbledon dilemma: is it better to have the best serve and extremely reliable volleys, or to return consistently well and have excellent passing shots in order to win at Wimbledon? To put the issue in biological jargon: is the strategy of serve and volley fitter than the strategy of returns and passing shots? As the history of tennis shows, it seems naive to give a specific answer: after all, Borg won 5 times in a row using the second strategy, while McEnroe won 4 times adopting the first strategy. Furthermore, many other players have won using mixed strategies. In this case I think we can talk of different "optimal" strategies, of "multiple peaks". In nature the situation is sometimes similar. Given a particular environment we can have different optima. Towards which one natural selection will tend depends on the starting conditions, and on the existing architectural constraints.

The last argument I present to rebut the optimisation hypothesis refers again to the nature of the environment. Even if it is granted that the cognitive strategies that constitute the science-forming capacity are genetically encoded and adaptations, there is no assurance that this constitutes the optimal cognitive system for science, given that a cognitive system optimal for hunter gatherers is not per se optimal for epistemic purposes. With this I am, again, neither disputing that adaptationist thinking is not a heuristically useful way of thinking about design problems, nor that it is not applicable to epistemology. Indeed, I believe that even without "absolute" optimality there might be rationality of some kind (i.e. minimal rationality). But the optimality assumption, unless it is properly tested with rigorous optimality models, is just a vacuous claim easy to challenge. I conclude that premises 1, 2 and 3 are idealised versions of untenable hypotheses that are hardly confirmed and universally accepted, and whose descriptive accuracy is under dispute.

Now I shall confine my treatment to premise (d), which provides the link between a putative optimality (granted for the sake of argument) and reliability, a link necessary in order to justify the science-forming capacity, to establish its rationality. The crucial question we need to ask is whether being conducive to survival and reproduction is equivalent to truth-conduciveness? According to Stich, premise (d) is based on uncritically assuming a link between optimality and reliability. In order to show this, Stich (1990 chapter 3) severs this link by contending that natural selection does not always favour reliable inferential strategies. In fact, given two genetically encoded inferential strategies differing in reliability, it is possible that the less reliable is fitter than the more reliable, and, hence, that natural selection will select for the less reliable one: "Often it is more adaptive to be safe than sorry" (Stich 1985 p. 123). Stich's "better safe than sorry argument" can be illustrated as follows. From a purely economical point of view it is easy to construct a thought experiment showing that a cognitive strategy so designed as to be extremely reliable is less fit. Natural selection would not select for a strategy so uneconomical, like, for example, that of deducing all the consequences from a certain belief (Sober 1981). Logical omniscience would certainly be an optimal, reliable and rational strategy, but it would not be the fittest, because it would drain the energy resources of the organism endowed with it (the economical costs in terms of memory space occupied, time spent in information processing etc., would certainly outweigh its theoretical benefits), at the expense of less theoretical but more practical inferential strategies.

Furthermore, from an inferential point of view, it is quite possible for a strategy that is less risky, more cautious, but less reliable than an alternative, to be the fitter one. Consider two strategies for detecting poisonous food. One generates beliefs on the basis of a minimum amount of evidence: as soon as the red colour of a fruit is observed, one avoids that food. The other is more reliable: just the observation of the colour is insufficient for generating avoidance, but some more information about the fruit (e.g. skin type, plant type, smell etc.) is required to generate avoidance. Which one of the two strategies is favoured by natural selection? The answer is that it depends on the nature of the environment. If in the environment 99% of fruits are red then a more discriminatory strategy like the latter seems to me both more reliable and more adaptive. However, if many red items do not constitute the food present in the environment, then we might have an adaptive advantage for the former and less reliable strategy. Furthermore, the first strategy is less risky because there is no possibility of generating false positive beliefs (inferring that the fruit is good when it is poisonous), while there is greater risk for generating false negatives (the fruit is bad when it is not poisonous). The reverse is true in the case of the second strategy. It is possible to argue that for certain kinds of cognitive tasks, given a certain environment, false positives are more dangerous than false negatives: given abundance of non-red food in the environment, then strategy one will again approximate less the theoretical optimum than the more reliable one, but will nevertheless be selected for because fitter.

The moral of the argument is that the concept of "theoretical optimality" as an environment-independent notion, as a mark of objective rationality, does not make much sense, because it is divorced from the notion of fitness. Reliability and optimality are therefore separated. Optimality itself cannot be guaranteed by natural selection. Not only it is not safe to presume that people adopt optimal (in an absolutistic sense of the term) cognitive strategies, but, furthermore, it should not be inferred that strategies that are more reliable have been selected for.

Stephens (2001) has questioned Stich's argument by pointing out that, even though Stich might be right in arguing that reliable cognitive strategies are not "generally" or "on average" selected for in the long run, in some circumstances natural selection will favour reliability.[56] What Stephens shows is that the general question Stich focuses upon has doubtful empirical and philosophical importance, and that it is better to consider under what environmental circumstances natural selection will favour reliability. The answer to this question is certainly context-dependent, since it depends on the nature of the environmental conditions and the nature of the cognitive task. Stephens' conclusion is that natural selection will favour reliable cognitive strategies when:

1) The cognitive task under consideration is "systematically" important, where systematicity means that the cognitive task is general enough as to occur in a wide variety of cases where having a reliable belief makes a lot of difference as far as biological fitness is concerned;

2) The environmental conditions are such that the organism has to consider and choose among several actions, each of which with important consequences as far as biological fitness is concerned.

Stephens' argument can be applied to Ruse's case. In this sense, what we need to establish is whether the cognitive strategies constituting the science-forming capacity help in solving systematically important cognitive tasks, and whether such cognitive tasks matter as far as biological fitness is concerned. If such cognitive strategies have, in brief, great survival value then it follows that natural selection will make them reliable. The point becomes whether Ruse is correct in assuming that "all" the strategies forming the capacity are so important as far as "biological" fitness maximisation is concerned (cf. sections 2.3. and 2.5). I believe that, in one sense, Ruse's characterisation of the science-forming capacity is too vague: modus ponens, induction and simplicity are not enough to produce science, even though I might admit that they are necessary to perform systematically important cognitive tasks in a wide variety of environmental conditions. Second, it is doubtful that the cognitive strategies that govern scientific behaviour are the same that were useful for our ancestors. In this respect, their "biological" advantage is doubtful, even though their "cultural" advantage is not.

For the sake of argument, let us assume that premises (a) - (d) are at least partly correct, and that the cognitive strategies constituting the science forming capacity are adaptive, reliable, and that they refer to local optima: does the conclusion follow? Stich contends that it might not:

"An inference pattern which generally gets the right answer in a limited domain is applied outside that domain, often to problems without precedent during the vast stretches of human and pre-human history when our cognitive apparatus evolved. Indeed, it is disquieting to reflect on how vast a gap there likely is between the inferences that are important to modern science and society and those that were important to our prehistoric forebears."[57]

This means that the cognitive strategies Ruse identifies could be considered irrational in the scientific context, even though they were selected for because reliable in the ancestral context. Stich stresses a genuine problem, even though Ruse seems to evade it, because modus ponens, induction and simplicity might be domain-general and not context-specific cognitive strategies. This prima facie evasion from the problem, however, hides the real limit of Ruse's account,

[56] I have somehow changed Stephens' argument. While I focus on strategies, he focuses on propositions. But the point of my different illustration remains intact.

since the more we ascend the level of complexity of the cognitive strategy belonging to the putative science-forming capacity, the more context-specific and dependent it becomes, with many more risks as far as reliability, and, a fortiori, rationality. I conclude that, even though, as an evolutionary epistemologist, I have no grudge against the argument from natural selection, this argument cannot be applied to justify Ruse's EEM.

2.5 The case against EEM: concluding remarks

The tension at the heart of EEM is determined by the fact that science does not confer advantages measurable only in terms of biological fitness. What is puzzling about science is that its value transcends genetic influence and direct biological advantage. In fact, as in many cultural activities, there exists an inverse relationship between biological and socio-cultural success. This in itself is a conundrum for sociobiology, since it shows that the central theoretical tenet of this discipline, namely that all human behaviour is selected for because biologically fitness-maximising, is false (Vining 1986). Scientists do not reproduce more, and the success of hypotheses and ideas is not identifiable with their biological advantage. The only way to solve this tension is by treating science as a cultural phenomenon. But only EET can provide the basis to treat scientific evolution as a cultural activity.

In this chapter I have not claimed that science is not somehow biologically founded, but rather that an epistemology that is based on such assumption is epistemologically incomplete. One defect of EEM is its biological foundationalism, which has been extensively criticised. For instance, Chomsky asserts that biology cannot explain how we discovered quantum mechanics (cf. quote at the start of the chapter). What Chomsky points out is that, in order to explain how and why knowledge grows, EEM leaves open an explanatory gap. Giere has stressed that to argue that the practical and theoretical cognitive strategies (those useful, respectively, "in the cave" and in science) are innate, and that the latter are based on the former, does not

[57] S. Stich (1985) p. 127

completely explain the emergence of science. This is what Giere (1996) calls the "1492 problem". Ruse realises the deficiency of his approach:

"...no one claims that the whole of science....has direct and immediate advantage. The claim is that such science is rooted in adaptive advantage - it is built around the principles informed and constrained by the epigenetic rules. It is these latter which are the legacy of the Darwinian process. But nothing says that science cannot take off from there, going beyond instant biological needs..."[58]

EEM provides only part of the complete explanation for the discovery of quantum mechanics and of the emergence of science. In this sense, the science-forming capacity as characterised by Ruse is a necessary, but not sufficient, condition for science. But what is crucial about Ruse's EEM is the claim that reference to the science-forming capacity is epistemically primary. I contend this claim is trivial. The point of Ruse's argument is not simply that we have innate knowledge of some kind that helps scientists in their cognitive activities. This would be an extremely uncontroversial claim. What is controversial about Ruse's claim is that reference to the science-forming capacity is of the utmost importance to explain the extremely variable and creative behaviour of scientists, and the pattern of scientific evolution. Ruse also contends that the biological foundation of scientific methodology is of the utmost importance in explaining the success of science. I believe that this view is false for two main reasons. First, because the science-forming capacity posited by Ruse is perhaps not as important as other selective processes operating in science, processes that more proximately account for the evolution of ideas and hypotheses, for the cooperation and competition between scientists, for the emergence of standards of selection that have no biological advantage whatsoever. Second, because the characterisation of the science-forming capacity Ruse provides is weak and trivial. The investigative strategy I have followed in the chapter has been to assess whether the cognitive procedures belonging to the science-forming capacity are innate and adaptive. I have shown that there are even doubts concerning the adaptive nature of modus ponens. To the further question whether the set of epigenetic rules proposed by Ruse

[58] M. Ruse (1986) p.170

provides a substantial characterisation of the "method" of science, my answer has been negative. Ruse's epigenetic rules are relevant as far as cognition in general is concerned, but not as far as science specifically is concerned. Given that EEM is necessarily restricted to a purely genetically based framework to account for science, it does not take into consideration the alternative and equally a priori acceptable hypothesis that the "method" of science consists of cultural adaptations (cf. section 5.3). For these two basic reasons, EET seems to me a better model to explain science's success.

I am not disputing that EEM provides part of the explanation of science's success, and that the features of scientific method that are supposedly instantiated as epigentic rules have explanatory relevance as far as the complete explanation of the way science evolves is concerned. Rather what I am disputing is that such "ultimate" explanation is of major interest. Epistemology takes for granted that scientists must be "endowed" with some cognitive capacities. In this sense, EEM is hardly a complete and primary part of any epistemology. For this reason, I completely agree with Bradie, who claims:

"Indeed, there is a sense in which some version of the EEM program must be true if our current understanding of evolutionary processes is anywhere correct." [59]

There are two further features of Ruse's EEM that I deem unsatisfactory: its methodological individualism and its epistemic cognitivism. The methodological hypothesis at the core of EEM is that, since scientific method can be purely characterised in biological terms, then the success of science supervenes somehow on psychological processes. However, it is doubtful that we can fully explain critical thinking and cultural (i.e. scientific) phenomena by mere reference to psychological processes. I believe that psychological explanations of the success of science are in this sense additional to historical and sociological ones, that is, to those explanations that EET, with its reference to social selective processes, tries to seek. To see

[59] M. Bradie (1986) p.408.

scientific creation and selection in purely individualistic terms leaves out much about the nature of the scientific process (cf. section 5.1).

EET also rejects EEM's "epistemic cognitivism", that is, the metaphysical hypothesis concerning the nature of the epistemic standards according to which epistemic standards have a biological foundation. Epistemic standards are based on human nature, and biology becomes epistemologically the most important science, being the basic science of human nature. The basic defect of this view is that it denies on a priori grounds the possibility that we participate in the "construction" of the standards (cf. section 5.3).

In this chapter I have criticised the "kind" of adaptationism endorsed by EEM, not adaptationism itself. EET is a much broader adaptationist hypothesis, since it regards the analysis of those selective processes peculiar to science that more proximally cause epistemic adaptations and the evolution of science. Adaptationism remains the most powerful hypothesis in evolutionary epistemology. This is because selection is the major force guiding evolution. This basic assumption is shared by both EEM and EET in different forms.

Chapter 3: The Universality of the Variation-Selection Model.

"The mind is more than a device for generating the behaviors that biological selection has favored. It is the basis of a selection process of its own, defined by its own measures of fitness and heritability. Natural selection has given birth to a selection process that has floated free."
Elliott Sober, "Philosophy of Biology", p. 220

In chapter 2 I showed that EEM is not a satisfactory model of the evolution of science. In this chapter I shall show that EET is better suited to provide a satisfactory model. However, some preliminary issues regarding the general nature of the model endorsed by EET will be considered.

First of all, EET should be constructed as a phylogenetic theory. There is some confusion at the heart of the distinction between EEM and EET, which is partly intertwined with the ontogeny/phylogeny distinction. Ontogeny/development and phylogeny/evolution are two different processes. EEM and EET come in two forms. Ontogenetic EEM is concerned with the development of the cognitive apparatus of the individual human/scientist. Phylogenetic EEM is concerned with the "evolution" of the cognitive structures of a species. Ontogenetic EET is supposed to show that science's success is a function of the knowledge development of the individual scientist. In this way epistemology is reduced to intellectual biography since the growth of scientific knowledge depends on the intellectual development of the scientists. EET should be considered as a phylogenetic theory. What EET is concerned about is not the cognitive development of individual scientists, an issue that should be left to the biographer of science. If EET is to be a theory of the growth of scientific knowledge it must be concerned with the historical evolution of science, not with the cognitive development of individual scientists. EET should aim at being a theory of the intellectual evolution of a population, be they humans, or, more restrictively, scientists. EET is alternative to EEM in that it is a theory of cultural evolution, and it is alternative to ontogenetic approaches in that it is not committed to methodological individualism (cf. section 5.1). The task of my thesis is to consider whether phylogenetic EET is a viable, coherent and explanatory epistemology.

In this chapter I shall defend the thesis that EET must be based on a Darwinian and selective model of cultural evolution. EET should describe and explain how a specific social mechanism of scientific selection (itself a variant of a mechanism of cultural selection) operates in science. Models of cultural evolution are generally criticised because they are not committed to biological foundationalism, which is the thesis according to which biology is more basic and fundamental than the social sciences, that nature is more fundamental (at least epistemologically) than culture. EET is therefore liable to be criticised for the same reasons. In my opinion, this is an open question. As Sober points out:

"There is a vague idea about the relation of biology and culture that models of cultural evolution help lay to rest. This is the idea that the science of biology is 'deeper' than the social sciences, not just in the sense that it has developed further but in the sense that it investigates more important causes."[60]

In particular, I shall show that cultural selection should be thought of as being a selection process completely detached from natural selection, and, rather, as a process whose action can be more powerful than that of natural selection. If this is true, then EEM's tendency to explain cultural phenomena by appealing to a more ultimate biological explanation is unfounded for a very fundamental reason. Even if brains are at the heart of the issue and everything can be explained in cognitive terms, even if many features of brains are adaptations and everything can be explained ultimately in adaptive terms, even if natural selection has produced the brain:

"....it does not follow from this that the brain plays the role of a passive proximate mechanism, simply implementing whatever behaviors happen to confer a Darwinian advantage."[61]

---

[60] E. Sober (2000) p.219
[61] E. Sober ibid. p. 219

The fundamental question for EET is to account for the workings of the mechanism of cultural selection which is "floating free" (cf. Sober's quotation at start of the chapter), even though it has been somehow created by the action of natural selection. This point has been emphasised by Skagestad:

"The crucial question in evolutionary epistemology is the question of how evolution by natural selection was able to generate, in one biological species, a mode of evolution not operating through natural selection, and yet contributing to the survival of the species in question."[62]

Even though I believe that Skagestad is mistaken in generalising about the biological adaptiveness of culture, the general concern of the passage is clear. Sober and Skagestad (both, in their own ways, detractors of evolutionary epistemology) agree on the fact that the growth of culture is a phenomenon that cannot be accounted for in purely biological terms. Dawkins stresses the same point:

"Are there any good reasons for supposing our own species to be unique? I believe the answer is yes. Most of what is unusual about man can be summed up in one word: 'culture'. I use the word not in its snobbish sense, but as a scientist uses it. Cultural transmission is analogous to genetic transmission in that, although basically conservative, it can give rise to a form of evolution."[63]

In order to solve the "mysteries" of cultural and scientific evolution, of how natural selection has given rise to modes of evolution and selection processes that are independent of its workings, we now need to examine what kind of model EET should be based on.

3.1 Selective and non-selective models of cultural evolution

Whether the application of evolutionary thinking can be extended beyond biology is a contentious matter. However, many philosophers agree that the matter cannot be settled a

[62] Cf. M. Bradie (1986) p.438
[63] R. Dawkins (1989) p.189

priori. Some are optimistic, others less so. Among the optimists, some compare evolutionary thinking to a "universal acid", that is a comprehensive and original way to approach many mysteries and antinomies that challenge common sense and that constitute the most intractable philosophical puzzles with which Western thought has been struggling to come to terms since the Greeks (Dennett 1996 esp. chapter 3). Some talk about the prospects of a "universal selection theory", extending the notion of evolution by (natural) selection to areas only distantly related to biology (Cziko 1995). In both cases the evolutionary agenda is vastly enlarged. Whether such an agenda can be enlarged as to encompass epistemology is the main question of this chapter. I believe that epistemology cannot be considered so detached from the rest of the traditional philosophical disciplines as to be completely intractable from the evolutionary perspective. I take it to be evident that evolutionary thinking could in principle illuminate certain problems of epistemology. Does the application of evolutionary biology to epistemology imply that epistemology becomes an empirical discipline, or that evolutionary thinking can encompass the treatment of normative issues? I shall give my answer to this question in section 5.2.

The resistance to embracing the evolutionary perspective stems partly from the urge to see humans as a special product of evolution, as somehow autonomous from the laws of Darwinian evolution, or, more mystically, as beings capable to revert and subvert the "necessity" of the putative evolutionary laws. Challenges to the sufficiency of natural selection as the directional force in organic evolution are as familiar today as they were in the past. New interpretations of saltationism, of the hypothesis of directed mutation, hypotheses concerning the self-organising properties of matter, different kinds of environmental instructionism etc., are proposed on an almost daily basis. The challenge for Darwinists is to find a way to reconcile the "autonomy" of culture and the emergence of "mind" with, first of all, a materialist metaphysics and, second, with Darwinism as selectionism. Life, altruism, consciousness, creativity, freedom can all be seen as "emergent" natural phenomena. The basic challenge for materialism is how to explain emergence in purely naturalistic terms. As far as I am concerned, I do not think that the mysteriousness of these phenomena can justify the rejection of materialism as the only acceptable metaphysics, the belief that science cannot explain them in principle, that these phenomena are not "natural". In this fashion any reference to Bergson's *elan vital* and vitalistic processes seems to me unnecessary. EET is based on the search for a general theory of evolution that can somehow heal the apparent fracture existing between organic and cultural evolution. I believe there are no a priori reasons why Darwinism cannot be applied to the study of science. Dawkins (1989 p.191) stresses this point as well: "I am an enthusiastic Darwinian, but I think Darwinism is too big a theory to be confined to the narrow context of the gene." The problem is to find out to what extent the Darwinian framework can be applied. But before doing so, it must be noted that non-selectionist and non-Darwinian models of cultural and scientific evolution have been proposed. Their historical importance, however, does not imply that they have any significant epistemic value.

Historically, the most prominent pre-Darwinian model for understanding evolutionary phenomena was taken from embryology. In fact, before Darwin, the term "evolution" was synonymous with development. The embryological model is also called "transformational" (Lewontin 1982), and is based on the assumption that the system of reference (be it an organism, a species, a society, or science) "evolves" by following pre-determined stages that are immanent or endogenous. Examples of this model come from Lamarck, and many theories of history (e.g. Comte, Hegel, and arguably Marx). Change in this model comes from the "inside", while in the Darwinian one it is exogenous. As Lewontin points out, the transformational model as applied in biology shows the idea of the hegemony of the genes (the internal "stuff"). Genes set up a developmental programme, while the function of the external environment is merely that of triggering the programme into action (e.g. Chomsky's hypothesis about the development of the language faculty). The structure of the developmental programme is "insulated", that is, largely unaffected by environmental influences. In a different way, Comte suggested to see societies as developing from a pre-scientific to a scientific age. For him, evolution is a finalistic and progressive developmental process governed by the unfolding of immanent and necessary stages. It seems to me that Comte's is an unexplanatory hypothesis that needs much elaboration. Its elaboration, I believe, would require the postulation of evolutionary processes that select the cultural items at the basis of the scientific world-view because they are fitter, in a non-biological sense, than the pre-scientific alternatives. Science does not seem to be a developmental process governed by the "inherent" tendency towards epistemic progress. Given a certain amount of optimism, science might seem to "develop" through a series of stages, but the idea that its evolution is just an "unfolding" of endogenously constrained stages insulated by environmental influences seems at best an unrealistic metaphor. Scientific, as organic, evolution is better seen as an opportunistic process, as a contingent and perhaps non-demonstrably finalistic process, in which no necessary steps are involved. In any case, epistemic progress cannot be but the result of a selective process. Note also that such a model is ontogenetic, not phylogenetic. Therefore, a further difficulty of the model as applied to science is to understand which is the "system" that develops, and why: is it the whole of science as an institution (i.e. the scientific process), or the cognitive agent (e.g. the scientist)? These two hypotheses are very different (Downes 1999). The second might be defended, but it is epistemologically unsatisfactory, since it does not refer to the social aspect of the growth of knowledge, while the first seems to me to require a selective auxiliary assumption. In general, however, I doubt there can be a "developmental phylogeny", since evolution is not a pre-programmed, insulated and finalistic process (Sober 1993 p. 153). Mine is not a proper critical "argument" against the significance of the transformational model. I am not disputing that the model might provide a sound basis for developmental psychology, and therefore it might have relevance for some parts of biology. As Lewontin (1982 p. 154) remarks: "It is important to note that there is nothing more evolutionary about variational evolution [Darwinian] than about transformational evolution." In principle, the model can provide the basis for explaining how and why science evolves, but it seems to me that it needs much elaboration, where such elaboration consists mainly in supplementing the idea of unfolding with the idea of selection. In any case, the burden of proof is on Comte's intellectual heirs.

The transformational model is non-Darwinian and non-selectionist. Cultural drift is, instead, a Darwinian but non-selective model. It is Darwinian because there is more to evolutionary theory than the theory of natural selection, in the sense that selection is not the only force causing evolution. In fact, the idea that evolution is caused by (random genetic) drift is familiar in biology (Kimura 1983). But random cultural drift cannot explain the emergence of "puzzles of epistemic fit", that is, the existence of some kind of relationship between our knowledge and nature itself. As in biology the hypothesis leaves open the problem of adaptation (Dawkins 1980), so is the case with a theory of scientific drift. Drift does not provide an alternative hypothesis to selection in order to explain the existence of design. This is because drift is the antithesis of selection (Sober 1993 p. 113). Drift provides chance explanations, which are epistemically and informatively limited in the scientific case, since science is a selective process where scientists select ideas and hypotheses by appealing to epistemic standards of some sort. With this, I am not disputing that stochastic processes have a role in scientific evolution as they do in biological evolution. I do not doubt that many cultural and even scientific phenomena are the result of drift,[64] but scientific evolution is not governed by chance alone, it is not a random walk. Therefore, I believe that random cultural drift provides only an auxiliary explanatory device to scientific selection.

The idea of cultural drift can be easily misinterpreted. In this sense, it is important to point out that, appearances notwithstanding, radical relativisys do not propose drift hypotheses. Many relativist epistemologists have put a strong emphasis on peculiar kinds of social factors that however retain a selective nature, despite being non-objective or non-internal. In this sense, epistemic externalism and social constructivism are not, by themselves, drift hypotheses. Even the extremist view of the late Feyerabend, according to whom theory choice

[64] Why does certain behaviour become "cool" and "trendy"? Certainly contingency plays a role in these cases, especially when such behaviour does not seem to confer any kind of biological or cultural advantage in terms of

is just a matter of "taste", is not a drift hypothesis, being taste somehow based on a kind of non-random psychological preference. "Anything goes by taste" is not a good definition of drift, since judging by taste whether to accept a hypothesis is a selective process. The hypothesis of drift does not provide any clue as to what the nature of the standards of scientific selection are.

I conclude that transformational and drift models of scientific evolution cannot be used as a basis for EET. Evolutionary theory provides a more powerful theoretical framework on which to base EET. However, many models of evolution by selection can be proposed. Only the least biological can provide a basis for EET. In order to build a proper model of scientific evolution, a major theoretical issue regards the relationship between EET and evolutionary theory. In particular, what is at stake is whether we should consider EET as mirroring in detail evolutionary biology, or whether a more elastic approach should be endorsed. Usually the fact that some fundamental distinctions present in evolutionary biology are not mirrored exactly or at all in a hypothetical science of scientific evolution is considered to be an a priori reason of why science cannot be treated from an evolutionary perspective (cf. sections 3.4, 4.5, and 4.6). I think this point of view is mistaken, basically because science is governed by specific selective mechanisms.

The relationship of epistemology to evolutionary theory can be accounted for in two ways. The first way of thinking about it is the orthodox extension of Darwinian theory. This conservative approach is based on the assumption that we can completely account for biological and cultural phenomena by using the notions of biological fitness and genetic heritability. Sociobiology and evolutionary psychology endorse this approach. The exemplification of this approach in epistemology is EEM as modelled by Ruse. As a matter of fact, this model is insufficient even to account for all biological evolutionary phenomena. This is because genetic evolution is not the only kind of evolutionary process that Darwinian

individual or group benefits. It might also be wondered why evolutionism has become so trendy in epistemology. To biology must account for. As Dawkins (1989 p. 194) puts it: "We biologists have assimilated the idea of genetic evolution so deeply that we tend to forget that it is only one of the many possible kinds of evolution."

In order to illustrate the difference between the conservative and the "enlargement" approaches, I shall rely on a categorisation of models of evolution by selection, which distinguishes three models, varying in their theoretical tenets (Sober 1994c).

Model I might be the basic model treated in biology textbooks, where evolution is defined as "change in gene frequencies". However, evolution cannot be defined in such strict terms. This is because the evolution of behaviour, which is the basic problem of a theory of cultural evolution, is only significantly phenotypic. First of all, the causal efficacy of a phenotype is independent of the nature of its genetic structure. Furthermore, if by traits or phenotypes all behavioural characteristic of the organism (including beliefs) are comprehended, then it is confusing to claim that a belief has a genetic basis. According to model I, purely phenotypic variation would not be heritable or transmissible from organism to organism, but this is obviously false if some learning process is involved. Without reference to learning, model I is not a satisfactory account of biological evolution as behavioural evolution. However, it might be argued that processes merely leading to phenotypic changes are of no particular importance to the evolutionist (Mayr 1976 p. 241). This is surely a simplistic view of evolution.[65] The vices of model I are three. First, the source of variation is not only genetic (e.g. mutation and recombination), but also ontogenetic (e.g. acquisition of new behavioural habits and/or beliefs during the lifetime of the organism). Second, purely phenotypic variation in biological fitness (like ontogenetic adaptations) exists (Magnus 1998); it can be heritable and transmitted to other generations (also horizontally) by means of some kind of learning process, and can therefore evolve (Boyd & Richerson 1985). Third, the final step of the evolutionary process is not necessarily genetic encoding.

---

this I reply that it is because cultural advantages will ensue.

[65] I am not claiming that this view is endorsed by Mayr, who uses the argument to make a point about "genetic" evolution.

In order to illustrate these vices, let us take an example. Consider the trait "starting a fire", which we might assume has been discovered by chance by some member of the Homo Erectus species (Leakey 1994). This discovery determined many morphological and physiological modifications. Fire means, among other things, cooking food and change of diet, and, therefore, it seems to me likely that such a behavioural set of changes triggered anatomic ones: modifications of the dental structure, of the digestive system etc. The hypothesis that ontogenetic adaptations do not affect structure is untenable. Genes are epistemologically irrelevant in the evolution of the fire trait for two reasons. First, the trait is heritable by learning. Second, there is no genetic encoding for the trait. Humans arguably are not born with the innate ability to start a fire, that is, there are no "genes" for the ability of starting fires. But the phenotype spread in the human populations nonetheless. How can such a remarkable phenomenon be explained? The only alternative is to posit the existence of a mechanism of transmission that is non-genetic. As a consequence, it would therefore be more appropriate to define evolution by selection as including change in phenotypic frequencies in a population due to non-genetic transmission of behavioural modifications.

Model I must be enriched in many ways to account for the evolution of behaviour, culture and science. We can assess now another less orthodox model of evolution by natural selection. Model II is based on non-genetic inheritance and biological fitness. Phenotypes are learned and are biologically advantageous. For these reasons they can evolve by selection and spread in the population.

Model II differs from I because of the characterisation of the notion of inheritance. Usually, the notion of heritability has a genetic connotation. For this reason, as Lewontin (1982) points out, such connotation should be substituted by the concept of "resemblance" between parents and offspring, which makes the notion of heritability broader and more appropriate. It is therefore sufficient for offspring to resemble their parents to have a condition of evolution by natural selection. Non-genetically based phenotypic resemblance is sufficient, as when offspring acquire new habits, behaviours or beliefs from their parents (for example, by imitating them). However, the concept of resemblance does not merely describe a relationship between parents and offspring, but also between genetically unrelated organisms. This notion is more appropriate if our aim is to describe cultural phenomena where transmission of information is not solely vertical, but also horizontal (Cavalli-Sforza & Feldman 1981). In this sense, model II is a more complete model of cultural evolution. EEM should be based on model II, though even this move is insufficient to save the epistemic primacy of the programme.

Another difference between models I and II is that in models of kind II genetic encoding is not necessarily the finishing point of the evolutionary process. Model II challenges the assumptions that genetic mutation is invariably the source of variation, and rejects the assumption that assimilation or encoding of the trait is the finishing point of the evolutionary process. When learning mechanisms are operative, assimilation can be bypassed. This is particularly the case with social learning (Tomasello 1997). The notion of evolution that models of kind II allow is much broader than that allowed by models of kind I (i.e. change in gene frequencies). Evolution is the vertical or horizontal transmission of genotypic or phenotypic variation in biological fitness. Consider yet again the fire trait. For the trait to evolve and spread in the population it is potentially sufficient that only one member of the population knows the skill; what we also require is the existence of a certain amount of cooperation, and the possibility of transmitting the skill non-genetically. That is, it is not essential for all members of the population to acquire the skill for the trait to evolve. This explains why there is no genetic encoding for this trait. Evolution by social learning prevents the cumulativeness of genetic evolution, because learning is "cheaper" than genetic encoding. However, even model II is insufficient as a basis for a theory of scientific evolution. This is because the notion of fitness in models of kind I and II is still biological, while science is not adaptive in any biological sense. Model III of evolution by natural selection is based on non-genetic inheritance and non-biological advantage. Model III alone can provide the basis for EET.

3.2. Model III and its significance for EET

Model III is strictly non-biological because it is based on a non-biological definition of fitness. In model III-based EET the reason for the differential survival and replication of the conceptual units of scientific evolution is not biological. There is a sense in which model III does not propose properly "biological" explanations (Sober 2000 section 7.5). It might therefore be concluded that the notion of "natural" selection is used purely as a "metaphor" in the model. I contend that science does not evolve only metaphorically, but that its evolution is governed by specific selective processes whose workings can be accounted for in Darwinian terms. Even though the explanations of model III are not biological, they are still "selectionist" explanations.

EET starts from the assumption that social learning is an operative mechanism of informational inheritance, and that it is a product of Darwinian evolution (in the sense that some cognitive mechanisms enable humans to engage in such a kind of learning). Model III is a Darwinian model based on the (cultural) selection of certain entities determined by their (cultural) advantage. Cultural fitness maximisation is about "having students", not "having babies" (Sober 2000 p. 215). The notion of evolution by selection that flows from model III satisfies the general definition of heritable variation in fitness, but expands it to phenomena in which phenotypic variation is transmitted either vertically or horizontally because of its biological or cultural fitness advantages. Model III respects the variability and heritability requirements, already satisfied by models of kind II. What model III adds is a new conception of fitness, which distinguishes it from models of kind II. While model II was still based on the very idea of natural selection, in model III we need to define selection in different terms. For this reason, model III is based on the postulation of a process of cultural selection. In the scientific case the major theoretical problem for a model-III-based theory of science would be that of explicating what the fitness, "transmission bias" or "attractiveness" of certain ideas, concepts and hypotheses amounts to. The problem concerns cultural (e.g. scientific) fitness.

Model III theorists must characterise both the nature of the selective process involved and the alternative notion of non-biological fitness.

As far as the notion of non-biological fitness is concerned, one difficulty in characterising it is that there are no obvious cultural analogues to viability and fertility. This is because, while the goal of living organisms is assumed to be that of surviving and maximising reproduction, in the case of beings capable of culture, the process of cultural selection has given rise to the emergence of different existential goals. Humans, in particular, have the power to change their aims, contravening the assumption supposedly valid for all other organisms, namely of striving for reproduction maximisation. For this reason, the analogy between biology and culture seems to break down concerning the notion of fitness involved. While in biology organisms' behaviour is understood in terms of biological fitness maximisation, in culture, given the potential multiplicity of aims and goals towards which organisms' behaviour is directed, we face the problem of understanding what are the units of selection, of which entities benefit from cultural selection: is it the individual, the group, or the phenotype itself? Of course the units of selection issue is a problem in biology as well, but at the cultural level the situation might be more complex. In the case of a regimented activity like science, space for teleological autonomy is more limited (science is governed by the pursuit of a limited set of aims).

The further issue that needs to be considered regards the ontological status of cultural selection. In model III cultural selection is an independent process of natural selection. This is the basic tenet of a theory of culture. But, is the postulation of a process of cultural selection evidentially sound? One of the reasons it is sound to posit an alternative process to natural selection is that some cultural phenomena can be seen as consisting of a clash between two selection forces. For instance, Cavalli-Sforza and Feldman (1981) propose to consider the switch in reproductive preference in nineteenth century Italy as a case of natural and cultural selection acting in different directions. The trait under discussion concerns the fact that Italian women passed from having, on average, five children, to having two children in the space of few generations. They explain the switch in preference by postulating the existence of a mechanism of cultural selection that acts as to favour the two-babies trait, which is biologically disadvantageous, but culturally fitter. This action overcomes the action of natural selection, which favours the clearly biologically advantageous trait of having five babies. In this case we have change that is purely phenotypic; furthermore, cultural and biological advantage do not coincide; we have evolution nonetheless, that is, change in the frequency of some phenotype in the population. The postulation of alternative processes to natural selection is not, therefore, condemnable in principle. The issue is rather whether cultural selective processes can be seen as relevant as far as human cognitive evolution is concerned. The debate between sociobiologists and cultural selection theorists concerns the magnitude of the selective forces shaping human cognitive evolution. According to sociobiologists, natural selection is the most important and most powerful force of evolution (in fact they use models of kind I and II). In case of clash with the hypothetical force of cultural selection, natural selection will win in the long run. Cultural selection theorists have different ideas: cultural selection can be stronger, even in the long run. Sociobiologists tend to treat cases of clash as instances of maladaptations: modern humans are maladapted to live in the contemporary environment and hence they do not act as maximisers of biological fitness. But they tend to believe that this is just a short-term accident. When presented with compelling evidence regarding the inverse relationship between fertility maximisation and social status among western humans (Vining 1986), sociobiologists tend to accept the evidence as an instance of a temporary accident.[66] This argument is fallacious. This is because of one main reason, concerning our "autonomy". Humans have the great (perhaps unique) advantage of being able to direct their behaviour towards the attainment of goals that are in principle completely unconcerned with maximisation of viability and reproduction. It is just in this sense that humans are "autonomous".[67] The whole issue is to clarify the nature of this autonomy in a

[66] cf. Symons' reply to Vining's paper.
[67] I do not pretend here to give an answer to the problem of free will, but just to point out that this axiological (i.e. goal or aim related) autonomy is the basis of freedom. I do not even try to explain how such autonomy is possible

way that is compatible with the general metaphysical framework of materialism and naturalism. Applied to cultural phenomena, Darwinism provides such a materialistic and naturalistic framework. Furthermore, what is clear is that without reference to the goals of scientists and science the most interesting epistemological problems remain untouched, and EET cannot fulfil its epistemological aims.

The significance of model III for EET is obvious. The model provides the means for understanding the evolution of certain phenotypes in spite of their Darwinian disutility, by solving the main theoretical problem of the EEM approach (cf. section 2.5). In fact, even though science is not adaptive in a biological sense, it is certainly adaptive from a more general cultural perspective. Model III explains why humans and scientists are multi-teleological creatures and cultural beings. The notion of scientific adaptiveness must however be clarified. It is certainly true that science can be seen from an evolutionary, Darwinian, selective perspective, but from this it does not follow that EET provides interesting explanations and normative recommendations concerning the scientific process. In this thesis I shall show that such explanations and recommendations do indeed follow.

Returning to the characterisation of the notion of fitness, I have said that the biological concept is clear enough, while the cultural one is not. One of the reasons for such vagueness is that cultural fitness can be measured in different ways, all of them highly contextual. For instance, Cavalli-Sforza & Feldman's explanation of why the two-babies phenotype is culturally fitter then the five-babies would be radically different from the explanation of why the notion of gradualism spread in the population of biologists in the last century. The concept of cultural fitness is not analogous to the concept of biological fitness, basically because culture is "more complicated" than biology. What I mean is that it is assumed that the "purpose" of biological organisms is that of surviving and reproducing, while such an assumption is impossible given the multi-teleological nature of human behaviour. The notion

(given determinism, or indeterminism), or how it originated. My suspicion is that freedom is an emergent of cultural fitness might be defined generally, but I shall be concerned with its scientific exemplification. Many items of scientific knowledge do not have any (direct?) biological advantage, but could enhance the prospects of the individual or the group in coping with the challenges of the cultural-social environment. Since scientific knowledge does not increase scientists' chances of survival or reproduction, a specification of its cultural fitness value must be proposed. The traditional point of view is that scientists are "truth" seekers or maximisers. But truth maximisation is not epistemologically on a par with reproduction maximisation, since "truth" has many meanings. For instance, Popper defined fitness in terms of verisimilitude, while according to Hull the scientific analogue of fitness concerns not primarily hypotheses, but scientists. Hull's notion of scientific fitness is conceptual fitness, which consists in scientists' efforts to maximise the transmission of their ideas to the population of scientists. In a word, scientists act as to maximise their "intellectual" reproduction. Hull's notion highlights the fact that EET is not necessarily limited to the explanation of the fitness of "ideas". The conception of conceptual fitness should in principle pertain to many different entities, as the notion of fitness in biology is not only relative to traits, but also to genes, organisms, populations etc.

One serious problem affecting EET as a model-III-based theory of scientific evolution is the following. EET would describe the evolution of scientific ideas but not the "causes" of why they evolve. That is, EET would describe the "consequences" of some ideas being fitter than others, but not the "causes" that render specific ideas fitter than others (Sober 2000 section 7.5). Therefore, EET as a model of scientific evolution would be uninteresting, given that it would only describe that some scientific ideas spread by explaining that they are fitter (i.e. epistemic-conducive). "Fitness" in this sense becomes a vacuous notion unless it is supplemented with an explanation that couples it with the notion of epistemic conduciveness. And this latter notion can be framed, for instance, either in internalist or externalist fashion. As Sober puts it:

---

evolutionary property.

"Evolutionary models of scientific change inevitably lead back to these standard problems about *why* scientific ideas change. It seems harmless to agree that fitter theories spread; the question is what makes a theory fitter."[68]

EET has open two ways to fill the explanatory gap. EET could be committed to provide an internalist definition of ideas' fitness, by preserving the normative dimension of epistemology: ideas are fitter basically because evidentially supported. In this way the evolutionary approach to epistemology would lose some of its bite, since what epistemologists find interesting is to focus on the internal criteria that govern evidential support. The evolutionary story would be a purely academic exercise, possibly providing some historical information, but nothing extremely interesting. Otherwise, EET could pursue an externalist solution to the puzzle: ideas are fitter if they are selected because better suited to solve the "environmental" (i.e. intellectual) problems faced by scientists. A new series of external causal factors is here taken into account to explain what scientific fitness amounts to, and in particular the role that the specific social selective processes characterizing the scientific process have in causing the spreading of particular ideas and hypotheses in the population of scientists. The evolutionary externalist position should however be distinguished from other and more relativist externalist positions (e.g. social constructivism). Popper and Hull (chapter 4) respectively epitomise the internalist and externalist approach to evolutionary epistemology.

Model III is the best framework for EET. It has the capability of explaining how humans are multi-teleological beings capable of culture. I shall now move on to illustrate the basics of Donald Campbell's EET model of scientific evolution. The particularity of Campbell's model III of scientific evolution is that it is based on the thesis of the universality of the variation-selection model of knowledge gain.

3.3 The thesis of the universality of the variation-selection model of adaptation

One basic assumption of Campbell's EET is that an account of scientific phenomena presumes an enlargement of the theoretical horizon of evolutionary theory, however still in line with the thesis of the universality of the Darwinian paradigm. In this sense, EET is based on the thesis of the universality of the variation-selection model of the emergence of adaptation. This thesis has been proposed by many theorists in different forms.[69] Selection processes are, according to "universalists", ubiquitous. In general, "universalists" believe that not only are Darwinian explanations better explanations of the presence of design than their alternatives, but that they are the only possible ones. In this sense, selectionism is not only a strategy of investigation, but the only intellectual paradigm enabling the explanation of emergent phenomena. For instance, contrary to instructionist explanations, selective explanations do not postulate the existence of guided variation induced by the environment (cf. section 3.4).

The peculiarity of selective processes is the kind of causation involved. Selective causation is dual: selection can direct evolution only if variation already exists. Both components are essential: without variation, selection has no "raw material" to select from, and change is impossible; without selection, evolution is a-directional. Arguably all major misunderstanding concerning the explanatory power of the variation-selection model (both in biology and epistemology) are ultimately due to the fact that causal duality is not thoroughly appreciated: evolution is not caused by one process, but by two separate but yet intimately interconnected ones. In this sense, there is no single ultimate cause of evolution. The dual perspective shows that the process is an interconnection of chance and creativity (Dobzhansky 1974).

---

[68] E. Sober (2000) p.218

[69] E.g. Dawkins (1999), Cziko (1995), Hull et al (2001c), Dennett (1996). Cziko, for instance, believes that the model is applicable not only to "life" phenomena, but also to artificial ones (e.g. genetic computing). Hull et al. argue that it is just an historical accident that the idea of selection was first developed in biology. Biological gene-based evolution is the paradigmatic instantiation of a selection process, but the explanatory power of the variation-selection model seems to be able to capture also the nature of immunological and behavioural evolution

In selection processes variants must be abundant. This is because selective processes must be wasteful in order to be creative. This means that necessarily some of the innumerable variants will not be taken into account by selection. Selective processes require iterativity, that is, the cyclical repetitivity of the processes of production of variation and cumulativity of selection. The directionality of evolutionary processes is basically explained in terms of the repeated cycles of replication, variation and environmental interaction (Hull et al. 2001c). Given self-replicating entities or entities replicating through self-replicating media, the wastefulness of variation, the cumulativity and retention of selection, and the iterativity of the cycles of variation-selection, the directional aspect of evolutionary processes is explained. The model provides a mechanical explanation for the emergence of adaptive complexity (Dennett 1996 pp. 48-52).

The variation selection model of adaptation has been criticised since Darwin first proposed it. Many evolutionary philosophers have struggled with the intellectual challenge brought about by Darwin. They thought that with its reliance on natural selection as the sole force moulding adaptive responses to the challenges of the environment, Darwin's theory was incomplete and could not therefore provide the framework for a comprehensive evolutionary philosophy. For this reason, Darwin's adaptive hypothesis had to be supplemented with further hypotheses. Criticisms of the Darwinian variation-selection model are varied. The most common one, which is applicable both to biology and epistemology, points to the fact that it is impossible to imagine that non-directed variations can produce complex structures. Another common criticism is that the variation-selection model describes a chance process, and that evolution is as a consequence unpredictable. Others point out that the production of variation is constrained, failing to realise that the notion of constraint is compatible and required by Darwinism (cf. section 3.4). Others claim that the variation-selection model does not explain the origin of replication, attributing to it an explanatory role it cannot play (Monod 1971 chapter 2). I believe that all these criticisms are misguided.

Even though I am agnostic about the range of application of the variation-selection model (its supposed phenomenal universality), and about the hypothesis that the model provides the best possible, or even the unique, explanation of any kind of emergent phenomena of functional order, I shall defend the thesis that the variation-selection model is a good explanatory model for epistemology (cf. section 5.2.2). This is because, in my opinion, the hypothesis of the universality of variation-selection is a sound working-hypothesis that has been repeatedly confirmed and not yet falsified (cf. section 3.4). Having said this, the application of the variation-selection model to science requires adjustments. In particular, intellectual variation is not as blind or random as genetic mutation. In this respect, I shall defend the thesis that the sense in which intellectual variation is directed is both Darwinian and fully accounted for by variation-selection models (cf. sections 3.4 and 4.5). Furthermore, the scientific evolutionary process is not mindless and ateleological like the biological one is. Applied to science, the variation-selection model must account for the "intentional" selection of variation. I shall therefore defend the thesis that, even if scientific and natural selection are two distinct and partly different processes, intentionality and selection are not alternative explanatory devices (cf. section 4.6). Furthermore, EET does not reject a priori saltational hypotheses. In all these senses, despite there being a number of disanalogies between biological and scientific evolution, in my opinion such disanalogies are not reasons to reject the validity of the thesis of the universality of the variation-selection model, nor the application of Darwinism to epistemology. That is, despite the disanalogies, EET remains a viable research programme. The fundamental reason why biological and scientific evolution are disanalogous pertains to the existence of peculiar social mechanisms of scientific selection that, despite being, in origin, products of variation-selection (i.e. some kind of selection process), have features that are distinct from "natural" selection. EET's task is to identify such peculiar features, and to describe and explain how the selective process of science is teleological (i.e. cognitive conducive).

EET's stance can be criticised for many reasons. For instance, Amundson claims that:

"As we proceed 'higher' on the hierarchical stages from [biological] evolution to psychology to social or scientific development, we should expect the explanatory force of naturalistic selection to gradually disappear. This is an ironic result of selection's own success at lower levels. The variation at a given level has been reduced, and in some cases directed, by selective processes at each lower level. At higher levels, the sorting mechanisms involve the purposive and insightful choice-making abilities that seem to have evolved in many organisms (including, for example, scientists). So both the source of the variation and the mechanism of sorting have become increasingly directed the farther one moves above the phylogenetic level."[70]

From my perspective (i.e. that according to which the disanalogies between biological and scientific evolution are not epistemologically important) this critical passage is completely in line with the general tenets of EET. The only difference concerns the attitude: while Amundson is pessimistic about the prospects of EET, I am not. And this is because what EET tries to identify are not scientific "analogues" of natural selection, but rather processes that have a selective function in science. In this sense, any strong reading of the analogy between natural and scientific selection is bound to be detrimental to EET's epistemological inquiry (cf. section 4.7). Even though analogies might have heuristic value, the scientific selective process should be analysed from an a-biological perspective. This point of view is endorsed by Campbell (1974a p. 434), who points out that "the valuable core for a theory of science is not the biological analogy to evolution per se, but a more abstract selection theory." In the same vein, Campbell (1974a p. 412) argues that evolution in its different forms (biological, cultural, scientific) is a knowledge process. The basic knowledge process involved is trial and error, or, in Campbell's terminology "unjustified variation and selective retention". The thesis of the universality of selection means that selection acts on many different entities (genes, phenotypes, ideas, conjectures etc.) at many levels of biological organisation (organic, cultural, social, scientific etc.).

EET should focus upon the peculiarities of the scientific selection process. In this respect, Campbell thinks of EET as mainly a social epistemology, which starts from an analysis of science that takes into consideration the "vehicle" of acquisition, distribution and validation of

[70] R. Amundson (1989) p. 430

knowledge. In this sense, EET is a descriptive theory of knowledge, where items of knowledge are not treated by abstracting from their physical realisation, but by taking into consideration the social or "vehicular" aspect. Traditional epistemologies generally reject this outlook, since a descriptive epistemology has no relevance as far as the validation of knowledge is concerned. This is not how Campbell thinks. What he shows is that an analysis of the vehicle is not only descriptively and explanatorily informative, but that it is also relevant as far as the context of justification is concerned. What is most illuminating in Campbell's approach is the emphasis on the distinction between those social selective processes that act as mere vehicle maintenance requirement, and those that promote epistemic fit maximisation (cf. sections 4.8 and 5.2.3).

Campbell (1974a p. 436) refers to the thesis of the universality of variation-selection as the "dogma" of his approach. Campbell's universality thesis is better seen as a working hypothesis, as a revisable and independently testable proposal with a metaphysical flavour. Campbell (1974b p. 142) argues that it is an "analytic" truth that in going beyond what is already known one cannot go but blindly (without foresight). Only some kind of guessing or "fumbling in the dark" can increase our knowledge. In biological evolution selection can act "creatively" (by resulting in adaptive changes) only if the raw material provided by genetic or phenotypic variation is already innovative (paradoxically, mainly in the forms of replication errors). The same is true in science. In this respect, he claims, there is no alternative to the variation-selection model, which becomes the only possible method of "inductive gain". What is peculiar about Campbell's views is that he does not show that the variation-selection method is the best available (or even the only possible) epistemic method. He also does not show that the method he proposes is successful. This is because the method cannot be justified (cf. section 5.2.2). Campbell gives up the traditional epistemological problem of justification tout court. It is in this sense that Campbell speaks of the variation-selection model as the "dogma" of his approach.

Given the importance of Campbell's views for my thesis, I shall constantly return to them in the following chapters, both for assessing Hull's views, and also in order to construct a viable

EET model. Campbell calls his epistemology "descriptive", but is eager to retain a normative role for such an empirical science of science. In particular, EET tries to solve epistemological problems from "within the framework of contingent knowledge, and by assuming such knowledge" (Campbell 1974b p. 141).

In the following section I shall present an illustration of how the variation-selection model is applied to the scientific case. I believe that such illustration provides a prima facie confirmation of the model of the emergence of knowledge endorsed by EET, while at the same time showing that arguments (particularly those based on the disanalogies between biological and scientific evolution) proposed to reject the epistemological validity of the model are not easily forthcoming.[71]

3.4 Vicariousness and the nature of variation

One of the supposed fundamental disanalogies between biological and scientific evolution concerns the different nature of variation: while biological variation is random, scientific or intellectual variation is directed, guided, Lamarckian (cf. section 4.5). Campbell's (and generally EET's) ambition of extending the thesis of the universality of the variation-selection model to all evolutionary phenomena, comprehending the evolution of science, would be dashed if such disanalogy (which can be labelled "Lamarckian challenge") is fundamental, given that the explanation of the success of science would be due to the directionality of intellectual variation, that is, to the process that governs the emergence of intellectual variation rather than to any selective process. However, I believe that Campbell proves that intellectual variation is produced in a way that is accountable in selective terms. Therefore the variation-selection model is saved.

[71] After all, EET's role is not only that of finding analogies between biological and scientific evolution, but also that of finding disanalogies between two obviously distinct processes. In this sense, the basic metaphor at the basis of the

In order to assess Campbell's argument, we need some preliminary information. Campbell's EET is adaptationist. Adaptation for Campbell means product of any selection process. Campbell's model differs from EEM models because it is "strongly" adaptationist. Avoiding a biologistic adaptationism implies a broader adaptationist stance (and the rejection of the epistemological completeness and primacy of the EEM programme). Evolution at all its different levels is governed by selectors of varied nature. Some selectors substitute the basic process of natural selection, and are for this reason called vicarious. Vicarious selectors are past results of selection, past "inductive achievements" containing wisdom about the environment, already phylogenetically achieved (biological and cultural) adaptations. Other kinds of selectors are constraints of varied nature that can be either by-product, epiphenomena, spandrels of selection processes, or even structural constraints limiting the optimising role of natural and other vicarious selectors. All these selectors constitute a nested hierarchy, acting at different levels of biological organisation that share the same functional role, by acting as mechanisms for the control of variation, in order to limit the amount of variation that is produced and selectable, to reduce, "a priori", the exploration of tentative trials that need to be taken into account to find the adaptive response, or, rather, the solution to the environmental problem facing the organism (or the groups of organisms). This enlarged notion of selector plays a central role in EET.

Campbell puts particular emphasis on the vicarious selectors that govern scientific evolution. One general problem is to justify their epistemic role. One kind of possible justification amounts to trusting the biological and cultural selective processes that have given rise to such vicarious selectors. One argument of this kind has been seen in section 2.5 (i.e. the argument from natural selection). In the form used by Ruse, the argument was invalid. If this were the only kind of justificatory procedure possible, then we would end up with a general scepticism about our chances of achieving knowledge. After all, vicarious selectors could be maladaptive:

evolutionary programme in epistemology can function, when opportunely modified, as an heuristic tool to rethink our views about the scientific process.

"The vicarious selector is only approximately accurate, and has fringes of inappropriateness which can produce illusions....Even in ranges of optimal function, the built-in presumptions are approximate truths about past worlds and may no longer hold if the ecology has changed."[72]

Furthermore, they could not only be unjustifiable, but false. In this respect, Campbell argues that:

"Our only hope as competent knowers is if we be the heirs of a substantial set of well-winnowed presuppositions...."[73]

However, a case for epistemic optimism can be stated by pointing out that scientific selection enables humans to engage in a reassessment of their multifarious presuppositions (for example by testing them empirically). In the remainder of this section I shall not focus on the justificatory issue. I will rather define more properly Campbell's notion of vicariousness. After I have done so, we shall assess Campbell's argument to rebut the Lamarckian challenge.

To illustrate the concept of vicariousness, I shall give a general definition, and then provide a few examples. The general definition of vicariousness states that x is a vicarious selector because it is a product of some lower level or more basic selection process (requirement about historical origin), and because x substitutes the selective role the lower level process had before its origination (hypothesis about substitution).

Habits and instincts are all vicarious selectors because they are generally products of natural selection, and because they substitute the selective role of natural selection. For instance, an organism that, at the fast approaching of an object, will respond instinctively by closing its eyes, will avoid a likely injury. In this sense, the organism, by reacting immediately to the environmental challenge, will adopt an adaptive response that was tested in the environment of ancestral adaptation by means of the life or death winnowing more proper to natural selection. Such more basic trial and error more proper to natural selection is thus substituted by the instinct. According to Campbell, unproblematic cases of vicarious selectors are also

[72] D. Campbell (1974a) p. 424

perception mechanisms (e.g. vision), cognitive processes (e.g. memory), learning processes (e.g imitation), and also language and consciousness. Consciousness is a vicarious selector in the sense that it is a product of selection processes (natural selection is the obvious candidate, but cultural selective processes might have a role in explaining its emergence without falsifying the validity of the variation-selection model, satisfying, that is, the first condition of adequacy of the general definition mentioned above), and in the sense that conscious states represent internally to the organism (i.e. mind) states of affairs that might necessitate adaptive responses that cannot be represented by means of the existing vicarious selectors, and that substitute the role of such vicarious selectors and of natural selection.

Passing to science, candidates for the role of vicarious selectors are the Kantian categories, which Campbell describes as follows:

".....highly edited, much tested presumptions, 'validated' only as scientific truth is validated, synthetic a posteriori from the point of view of species-history, synthetic and in several way a priori (but not in terms of necessary validity) from the point of view of individual organisms."[74]

The Kantian categories are in this sense interpreted as adaptations of some kind that substitute the selective role played by lower level adaptations. The requirement about origin is satisfied if it turns out that the categories are adaptations of some kind. The second condition of adequacy is satisfied if it is shown that the category substitutes a less advantageous (biologically or culturally), and more basic cognitive constraint. For instance, endorsing the belief that every effect has a cause substitutes the belief that some events have no causes (assuming, for the sake of argument, that this is a more basic belief). The crucial addition made by the EET perspective is that the concept of "category" can be extended so as to encompass both "biotic" (biological functions having survival value) and "sociotic" (cultural products) results of selection processes (Campbell 1974a p. 445).

---

[73] D. Campbell (1977) p. 444
[74] D. Campbell (1974a) p. 441

Other vicarious selectors are those social features of the social organisation of science that have an adaptive role, in the sense that their causal role is effective in furthering the institutional aims of science. For instance, conservative behaviour in science (i.e. that kind of "deviant" behaviour that scientists sometimes manifest by resisting putatively successful refutation attempts, that is, by interpreting apparently disconfirming evidence as not anomalous) is an adaptation in the sense that it is necessary to the effective working of the social system of science (i.e. without such kind of behaviour nobody would defend the hypothesis under test by trying to improve it, by developing it further etc.). Furthermore, conservative behaviour substitutes a more basic kind of vicarious behaviour, namely, the uncritical (e.g. unconditional) response to empirical disconfirmation.

It must be admitted that the more we ascend the level of biological or cultural organisation, the more problematic Campbell's notion becomes. The notion of vicariousness as applied to the scientific case is problematic because it is difficult to characterise how the two conditions of adequacy are satisfied in science. This is because the concept of adaptation in epistemology is not as specified as in biology. The concept of scientific adaptation is unclear until a more general theory of the functional organisation of science is forthcoming.

Generally, we can say that while the biological kind of vicarious selectors refer to individualistic features of scientists as organisms (e.g. vision, memory, ability to use a language etc.), the social kind refers to sociological features pertaining to the adaptive organisation of the social structure of science (e.g. institutional norms of science, etc.). Furthermore, vicarious selectors can be psychologically (vision, memory, language, consciousness etc.) or socially (mechanisms of cultural conformation, social methods of knowledge acquisition, socially vicarious learning processes, etc.) instantiated (Campbell 1974a p. 421). While the emergence of biological vicarious selectors is the fundamental problem of EEM, the fundamental problem of EET is to identify, describe, and explain the workings of the cultural vicarious selectors that render science a "socially distributed selective system".

It must be stressed that at the cognitive level, the concept of adaptation is not enough in order to characterise the nature of the multifarious cognitive constraints affecting the production of intellectual variants. But this is not a problem of Campbell's view, since the notion of vicarious selector is coupled with the notion of intellectual constraint (whose biological analogue is the notion of biological constraint on perfection or optimisation), which refers to all the cognitive biases affecting the production of knowledge (e.g. education, training, ideology, intellectual inheritance, previous knowledge, etc.). Intellectual constraints might therefore be instantiated as specific beliefs about the methodological and metaphysical aspects of the scientific endeavour.

Campbell's basic point is that the "economy of cognition" is based on the workings of a hierarchy of vicarious selectors and on a variety of intellectual constraints that ultimately act so as to reduce search space and amount of intellectual variation. I refer to such selective requirements as "preliminary selection". By means of this notion we can make sense of the claim that there is a link between the emergence of variation and the environmental situation in which the scientist finds herself. Campbell's claim is that the directionality of intellectual variation is caused by the action of preliminary selectors. I shall now illustrate and assess Campbell's argument.

Scientific evolution is said to be Lamarckian because intellectual variants arise because they would be useful for selective purposes. The analogy between biological (where the Lamarckian hypothesis has been falsified and rejected) and scientific evolution breaks down when it is claimed that the environmentally influenced ontogenetic responses of the scientists are directed at and beneficial to the solution of the problem they face. If the claim means that scientists produce variants with foresight, that scientists have an unspecified capacity to produce variants at will in directions that they already know will be successful, then Campbell rejects this claim. This is because scientists cannot anticipate whether the hypotheses they formulate will be selected. Scientists do not "know" with foresight what the selective process will lead to. This is the sense in which hypotheses are tentative. In order to explain this concept, Campbell claims that intellectual variation is "unjustified". However, critics argue that Campbell's notion of "unjustifiedness" is not properly explanatory. This is because in science there is a link between the occurrence of variation and adaptiveness, and therefore there remains a disanalogy between, for example, scientific variation and genetic mutation. What Campbell claims is that even in this sense the disanalogy between the two processes is not fundamental. In order to show this, Campbell refers to the largely misunderstood and neglected role of the notion of vicarious selection, which partly allows one to show in what respect scientists' hypotheses are not blind or random, but guided:

"Let me confess that I concede most of an opponent's facts from the very beginning [i.e. that hypotheses seem to arise as already adapted to the nature of the environmental challenge]. With my nested hierarchy of vicarious variation and selective retention processes, there is no need *for the observable products* at any stage to be blind."[75] (my italics)

A crucial point is that the formulated hypotheses (the "observable products" of the passage) appear as pre-selected because all the largely unconscious mental processing (cognitive and non cognitive) that led to formulation has no observable trace in the final stage. The other crucial point is that any evidence of pre-adaptation, pre-selection, coupling, directedness of variation, has to be partly explained by appealing to the notion of vicarious selection:

"It is not denied that if we observe variability (in problem-solving behaviour, for example) we may find the variations intelligent, preadapted, etc. Rather, if we find them thus, this is taken as evidence of already achieved adaptation, rather than a relevant explanation of further increases in adaptation."[76]

It might be contended that Campbell's point is weak, since he only shows that any evidence of coupling between intellectual variation and selection can be explained "away" by appealing

[75] D. Campbell (1974b) p. 152
[76] ibid. - p. 147

to the concept of already achieved adaptation (i.e. vicarious selection). Campbell challenges this criticism by claiming that:

"For this adaptive bias in variations is itself evidence of fit needing explaining. And *the only available explanation* (other than preordained harmony) is through some past variation and selective retention process."[77] (my italics)

Campbell's argument is based on the already familiar assumption that the variation-selection model of adaptation provides the best possible explanation when evidence of emergence appears, simply because it is the only acceptable explanation, given that providence and instruction are not viable alternatives. What Campbell wants to achieve is to shift the burden of proof onto those who do not accept selectionist explanations of this kind. If the critics can provide any alternative explanation, then EET has to consider yet again such a criticism. The notion of vicarious selection is crucial in Campbell's argument, but another crucial element concerns the influence of chance and accidental factors in hypothesis generation:

"...variations could be by chance and the process would still work. The source of the variations is irrelevant."[78]

Campbell argues that the directionality of variation is not even "essential" for the emergence of epistemic fit, since the fit between representation and represented can happen accidentally. This view does not imply that guided variation is not an endemic feature of scientific discovery, but rather that guided variation renders the process less wasteful and quicker. However, what is essential is just selection, not the source of variation, in order to explain science's success. That is, the source of variation is explanatorily irrelevant as far as the explanation of epistemic fit (i.e. science's success) is concerned.

I reconstruct Campbell's argument as follows. He makes four claims. The first is that variation is guided because it already involves (conscious and unconscious) preliminary selection. The

[77] ibid. - p. 150
[78] ibid. - p. 147

second is that science can in principle progress even without guided variation. The third is that the explanation of science's success depends on selection, but not on the nature of variation. The fourth and more general claim is that selection theory provides the only available explanation of creative processes. I shall assess and defend these claims in turn.

The first thesis Campbell defends is that the directionality of intellectual variation is caused either by the selective action of vicarious selection processes or by the selective role played by other kinds of intellectual constraints, namely those beliefs that scientists implicitly accept and that make up their particular perspective, their "hopefully largely coherent revisable system of beliefs of a multifarious nature". Vicarious selectors and intellectual constraints together constitute what I term preliminary selectors. EET allows that scientists do not start from scratch in the search for useful hypotheses. EET allows that already achieved "knowledge" plays a selective role in the scientific process. EET admits that scientists have already a preliminary idea about what constitutes a solution to a problem. EET also admits that scientists create hypotheses and ideas by trying to solve a specific problem, by having a specific goal in mind (in this sense there is a genuine difference with the production of genetic mutation), but, at the same time, this does not mean that such directionality goes beyond the constraints imposed by preliminary selection.

It might be argued that this explanation is not sufficient in order to make sense of the difference between scientific and biological variation, since the biological concept of randomness lacks an element that pertains to the scientific counterpart, namely that in science "wisdom of later variations is improved by the knowledge of the failure of the earlier ones" (Campbell 1974b p. 150). However, I contend that the notion of preliminary selection is sufficient to explain such a difference. As far as science is concerned, the variation-selection model explains what features of science make it possible to rely on the knowledge that some trials were either beneficial or non-beneficial. For example, at various stages in the history of evolution, vicarious processes were produced that rendered the process of design exploration less wasteful. This nested hierarchy of vicarious selectors substitute the trial and error processing of lower level selective processes by producing a tentative trial more quickly, for example as an instinctive ontogenetic response to close one's eyes when objects are dangerously approaching the retina. This is possible because the vicarious strategy contains an element of knowledge about past and unsuccessful trials. In this way, science improves epistemic fit more quickly than genetic mutation improves the fitness of biological populations, given the amount of vicarious selectors involved.

It might be argued that, still, there exists a basic difference between scientific discovery and genetic mutation. It is certainly true that genetic mutations are not more likely to be beneficial after *n* trials (where *n* is a very large number). This means that the occurrence of a beneficial (or deleterious) genetic mutation does not become more probable with time. The same, critics argue, cannot be said about intellectual variation. This is because there exists a second connotation of the concept of randomness as applied to science that is lacked by biological counterparts, a connotation that pertains to the fact that intellectual variation is more likely to be beneficial than completely unbiased variation (i.e. supposing genetic mutation is completely unbiased, which is not). One way to make sense of the peculiar nature of science is to claim that in science the possibility to come up with novel hypotheses is much more constrained than in genetic mutation.

To illustrate the point, compare two kinds of adaptive responses triggered by the emergence of an environmental stress. Suppose we compare a population of bacteria and a population of scientists, the first facing the presence of too much sugar in the environment, the second the problem of tackling the epidemic of a new virus. In both cases let us assume that the environmental stress increases the rate of variation. What I ask is whether the increase in variation rate will necessarily make adaptiveness more probable, and whether this phenomenon is Lamarckian in any sense. To the first question I answer positively. The population of bacteria will mutate more frequently, and as a consequence an organism with a

beneficial mutation will emerge more quickly. This organism will have more chances of reproducing and its novel and beneficial trait will eventually come to fixation if the environment remains stable. Scientists will similarly start to compete and cooperate in order to identify the virus, while public and private money will be spent to find and develop virus-resistant drugs. The process of intellectual variation will be speeded up, and as a consequence innovative treatments will be discovered more quickly. These treatments will eventually spread in the human population if the epidemics persist, and if the pharmaceutical companies do not prevent the access to the treatment for their usual reasons.
The second question is whether the process is Lamarckian. In the case of bacteria, the process would be Lamarckian if and only if the environmental stress made more probable the arising of a beneficial mutation than that of a neutral or deleterious mutation. What is rendered more probable in this sense is not such arising, but only the spreading of the beneficial mutation in the bacteria population, given its selective advantage. The process is not, therefore, Lamarckian. The same seems to me to be the case in the scientific case. The emergence or arising of a new treatment is not more probable than the arising of a bad or ineffective treatment. The only difference between the two cases is that scientists know what the problem they are trying to solve is, and that, unlike bacteria, scientists know what kind of solution they are looking for, what kind of information is relevant to find a better treatment than those available. But the process cannot be said to be Lamarckian just because we have such knowledge.

Ultimately, I think that to compare genetic mutation and scientific discovery gives rise to misunderstandings. This is because mutation is much less affected by preliminary selection than, for example, the evolution of behavioural responses. If selective processes constitute a hierarchy, as Campbell argues, we would expect the extent of preliminary selection (or already achieved knowledge) to affect the production of variants the more we ascend the hierarchy.

I contend that to call scientific evolution a Lamarckian process is partly a semantical issue. The directionality of intellectual variation is a fact, but this is not enough to call the process Lamarckian. This is because the selective model proposed by Campbell explains such directionality in terms of the various selective constraints that influence scientific discovery. The only meaningful way in which I suggest scientific evolution is a Lamarckian process is in the presence of foresight, which is not demonstrably involved in the process of discovery. The other sense in which scientific evolution is Lamarckian is that generated hypotheses are more likely to have beneficial rather than deleterious or neutral effects. However, this is partly false, and partly accountable in selective terms by appealing to the notion of preliminary selection. It is partly false because it must be stressed that the formulation of a hypothesis is just the final step of a very long selective process that is difficult to reconstruct. Without foresight nobody is able to decide whether a formulated hypothesis will have a beneficial effect on the solution of the problem at hand. This judgement can only be given a posteriori. If such a judgement is only based on the assessment of the final stage of the long process of hypothesis formulation, then the "appearance" of directionality seems to come from nowhere. But if the judgement is based on a fuller reconstruction of the process of scientific discovery by including all the false steps taken by scientists in the route towards formulation (i.e. all the unconscious and conscious cognitive and non cognitive strategies that they use in order to eliminate the potentially infinite array of hypotheses that could be beneficial in solving the problem), then directionality can be accounted for in selective terms. My view is that the presence of guided variation is just the final-stage manifestation of a profoundly wasteful process. Scientists come up with a staggering number of bad ideas, but at the time in which the hypothesis is formulated all these bad ideas have been selected against. Campbell also provides anecdotal evidence confirming the causal role of preliminary selection coming from studies in the psychology and sociology of discovery. The notion potentially explains why several people independently invented many major innovations in the history of science roughly at the same time (Campbell 1981 p. 490). This is because the (roughly) same selectors affected (roughly) in the same way the mental operations of the scientists.

The second claim Campbell makes is that science can in principle progress even in the absence of guided variation. If one wonders whether this hypothesis can be coherently sustained, then it must be shown that it is somehow confirmed by historical evidence. One way to show this is by appealing to episodes in the history of science in which accidentality has played a major role. Fleming discovered penicillin at St. Mary's Hospital in London by chance. Kekule' discovered the structure of the molecule of benzene when hallucinating. In all these cases, these scientists contributed to the growth of knowledge "accidentally". Campbell's thesis is a purely theoretical one, since "retention" is impossible to eliminate. Thus, Campbell's argument is not that the mind of scientists is just a chance-like mechanism. Selection theory makes no such claim, since preliminary selection is a basic and fundamental part of the creative process of discovery. For instance, Kekule's discovery was not purely accidental or chance-like, because some kind of preliminary selection was involved.

It could be argued that this model of hypotheses' generation might explain only part of the issue, but that in at least some cases (or perhaps the vast majority) such a model is not instantiated. For instance, it could be argued that the issue could only be solved by testing the hypothesis that "insightful", "purposeful" and "directed" hypotheses are much more likely to be successful than hypotheses that were generated by the interplay of chance and preliminary selection (Amundson 1989 p. 427). Let us assume that Campbell's argument is not convincing. I wonder what could be done in order to test the sceptics' hypothesis, and in order to convert them to Campbells' paradigm: should we ask scientists to describe their mental processes in order to show that creativity is actually a foresighted (in the sense of Lamarckian) venture? Can we trust such information, given that unconscious processes largely affect the psychology of discovery? Should we resort to phenomenology instead, or perhaps ask scientists to keep well awake, and force them to write introspective reports or scientific analogues of Joyce's "Ulysses"? The critics' hypothesis is actually untestable. It seems to me that EET's commitment to the thesis that the source of variation could be "unjustified" (which is not the same as chance-like) is usually resisted simply by rejecting a priori the possibility that scientific discovery can be partly random, blind, spontaneous, accidental, and partly the outcome of preliminary selection based on retention. Such a reaction might be understandable, but it is puzzling when it comes from those who accept the thesis that there is no logic of discovery. In fact, EET provides the sketch of a "logic" of discovery in the form of a model that can work as a basis to understand the supposedly unexplainable nature of creativity. I do not see anything remotely similar coming from the other side, which is content to remind us that "it cannot be true" that scientists' psychology is just a mixture of chance and preliminary selection. As a matter of fact, I believe that, on the one hand, whenever critics have to specify the nature of their arguments, they fail. On the other, when they manage to specify arguments supposedly showing the incompleteness of the variation-selection model, they produce bad arguments.

For instance, Amundson (1989 p. 429) claims that in order to explain the success of science "...the burden of explanation is carried by the remarkable and as-yet-unexplained abilities of humans to devise productively imaginative hypotheses..." I contend that appealing to an unspecified ability to come up with insightful hypotheses is hardly a recipe for understanding why science is successful, and hardly a way to increase epistemic knowledge. While recognising that there exists a selective side to the process of hypothesis formulation, Skagestad sees a disanalogy between scientific formulation and the action of natural selection. In fact, Skagestad argues, while natural selection always operates by selecting favourable variants, in science hypothesis generation is guided by pre-selection processes that exclude a wide range of guesses, "some of which might be true."[79] This argument is vitiated by the assumption that natural selection will necessarily "optimise", which amounts to a Panglossian thesis. Furthermore, Campbell would not object to the fact that some valid intellectual variants will not be taken into account because of our biases. It is rather a fact that "truth" can be lost on the way.

[79] Cf. Bradie (1986) p. 425

The critics of EET do not provide an alternative hypothesis about theory generation, and they do not contribute to understanding such a process. All we are left with is just a series of vacuous claims that seem to imply the presence of foresight in some guise or another.

I suggest that perhaps the only way to make sense of the appearance of "insight" is by relying on notions that are compatible with selection theory. One of these is error in replication, a notion that has central importance in biology. Error in memetic replication can be seen as a central strategy to introduce intellectual variation in science. Recombination is another such notion. Appearance of insight can also be explained by what can be described as "change in perspective" (i.e. cases in which a belief - or group of beliefs - does not function any more as a constraint on hypothesis generation). In all these cases selection theory makes no appeal to the involvement of foresight. Rather, the nature of the creative process is due to the interplay of accidentality and preliminary selection based on knowledge retention. Since "design exploration" is limited by the action of various selectors (with the consequence that scientists will consciously and unconsciously eliminate alternative hypotheses), the process of hypothesis generation will be far less wasteful, and hopefully quicker, than the analogous biological process.

Campbell's third claim is that the source of variation is irrelevant to explain the success of science. This means that selective explanations are complete, and that their explanatory force is retained despite the nature of intellectual variation. Critics contend that this claim would be disconfirmed if it were established that biases affecting the source of variation had an explanatory role. The strange thing about such criticism is that EET does not deny such a causal role. Campbell's model of creative thought shows that an explanation of the success of science is not complete without reference to such biases. Campbells' third claim only means that any explanation of the success of science must "start" from the reference to some selection process, and that whether an idea originally cropped up just by accident, error, luck or guess is irrelevant. Therefore, the power of selective explanations remains intact.

I believe that another reason why Campbell's model is easily misinterpreted is determined by the difficult articulation and specification of the notion of preliminary selection. This is because the notion of preliminary selection not only refers to adaptive features of both biological and cultural origin, but also to various kinds of belief, states of mind, and in a nutshell any item of information that has the power to decrease the amount of intellectual variation produced. Such a notion seems so malleable that it allows either to explain "away", or to find ad-hoc explanations of "genuine" cases of foresighted variation (if they ever happen). In fact, Campbell defines knowledge as a product of selection. Having said that, I still believe that what the variation-selection model offers is an explanation that is certainly incomplete, but that has no alternatives. It is incomplete since the reference to the vast array of biases affecting the process of hypothesis generation is impossible to articulate in complete detail. But it has no credible alternatives. Providence is not one. Lamarckian instruction comes in two forms: the first is foresight, but foresight is not demonstrably involved in scientific discovery; the second is directed variation, which is explained completely in selective terms. Taken seriously, Campbell's model is not in the least disconfirmed by available empirical evidence, since hypothesis generation can be accounted for in terms of the interplay of chance and preliminary selection. I therefore follow Campbell in challenging anyone to come up with an episode in science in which insight or creativity cannot be explained in these terms. The absence of disconfirming evidence and the lack of alternative models of scientific creativity suffice to vindicate Campbell's fourth claim, and to shift the burden of proof onto the critics of EET.

In this chapter we have seen how EET differs from EEM. The basic reason is that, while EEM is based on a model of evolution by selection that is too limited (i.e. model I or II), EET is based on a model of evolution that posits the existence of a peculiar selective process, and on a notion of fitness that is purely scientific and non-biological. In this chapter I have also introduced the nature of the variation-selection model and the basic tenets of Campbell's EET.

EET is based on the working assumption that the variation-selection model provides a useful and fruitful way to consider the nature of the scientific process (cf. section 5.2.2). The application of the model I provided was supposed to show both that arguments based on the existence of fundamental disanalogies between biological and scientific evolution are not easily forthcoming, and that the epistemic credentials of the variation-selection model are at least partially vindicated.

In the next chapter I shall assess the vices and virtues of Hull's EET model. Campbell's ideas and insights will play an important role in this assessment.

Chapter 4: Hull's Science of Science

"The mechanism that I propose rests fundamentally on the relations which exist in science between credit, use, support, and mutual testing. Science functions the way that it does because of its social organization....This mechanism is an instance of a selection process, but it is social, not biological."
David L. Hull, "Science as a Process", 1988, p. 281

Hull has contributed extensively to the literature in the philosophy of biology, particularly on the issues of the individuality of species and the units of selection. Hull's main contribution to epistemology has been that of elaborating an evolutionary view of science based on the action of a social selective mechanism. Hull's epistemological aim is to explain the growth of knowledge. The investigative strategy pursued is to look at the practice of science, by pinpointing the causal factors influencing scientists' behaviour, and how this behaviour tends to produce the growth of knowledge. In particular, Hull notices two crucial facts about how science is practised; first, that scientists tend to strive for reward and credit from the community of scientists (in Hull's terminology, maximise their conceptual inclusive fitness); second, that scientists tend to organise themselves in groups. The notion of scientific fitness and demic structure of science, Hull contends, are crucial in order to understand the evolution of science.

In this chapter I shall first give a brief presentation of what I consider to be the main epistemological tenets of Hull's EET. I shall then focus on his mechanism of scientific selection. I will then assess Hull's views, in overtly critical, but sympathetic, fashion. I shall then finish by considering the implications of Hull's EET.

4.1 A Theory of Socio-cultural Evolution

Categorising Hull's epistemology is not straightforward. In fact, Hull rejects the label "epistemology" for his theory:

"A common objection to evolutionary epistemology is that it is not 'epistemology'......I agree with the critics that a purely descriptive epistemology is 'epistemology' in name only.....I find a scientific theory of sociocultural evolution a vastly more significant goal than an evolutionary epistemology. A theory of sociocultural evolution, like any scientific theory, is more than merely descriptive, but any necessity it may have is nomic, not epistemic. My use of the appellation 'evolutionary epistemology'....should not be taken as an endorsement of any epistemological views whatsoever. If evolutionary epistemology were a genuine epistemological theory, I would not be in the least interested in it." [80]

Hull defines his EET as a "science of science". An empirical theory about scientific evolution can be categorised only unconventionally as an epistemology, but this is at least partly true of any naturalised epistemology. The issue is about the nature of naturalised, and, in particular, evolutionary epistemologies. The problem is about normativity: how can an empirical theory concerning the practice of science lead to the extrapolation of norms of action for scientists, and how is it possible to find out about necessary and/or sufficient conditions underpinning scientific success and evolution? As the passage above discloses, Hull is clear enough in pointing out that as far as the science of science is concerned, the sense of normativity concerned is nomic, not epistemic.

More specifically, Hull contends that the only alternative to the empirical evaluation of normative claims is the appeal to intuitions (Hull 1990 p. 80). And he rejects this alternative justificatory procedure because any intuitionist account is fatally flawed, being that our intuitions are subjected to empirical assessment, and because among alternative intuitionist theories of justification no criteria of choice are available apart from empirical ones. Of course, this line of reasoning is in line with a naturalistic and empirical approach to epistemology. Hull then prefers to consider a kind of reflective equilibrium as the only suitable justificatory strategy:

"Extensive knowledge of science should inform normative judgments about knowledge acquisition, while these judgments in turn should be evaluated by putting them into practice to see what happens. The best way to support the claim that scientists should not waste their time looking for laws of nature is to convince scientists and see what happens. If the result is an increased growth in our understanding of the natural world, then I for one would be convinced. I can think of no stronger support for normative claims."[81]

Hull also rebuts the criticism that there must be a fundamental distinction between normative and descriptive (i.e. empirical) claims about science, contending that:

"My account of science is intended to have both dimensions. I argue that science succeeds in realizing its goals more successfully than do other social institutions in large measure because of the social structure that it has evolved. In particular, the good of the individual scientists usually coincides with the institutional goals of science. Rarely are scientists asked to risk their individual careers for the good of science as such. The general story I tell has clear prescriptive implications, but as I see it, the only way for these prescriptions to be evaluated is for scientists themselves to put them into practice and see what happens. What should be the case does not always turn out to be the case."[82]

In brief, the normative claims Hull hopes to articulate consist in derivations from nomic generalisations. These generalisations hopefully capture "laws" about the scientific production of knowledge. They, that is, point to necessary conditions for the success of science. From these, recommendations can be derived. The problem of the justification and adequacy of these recommendations or normative claims is empirically solved. Evidence has a major role in testing these claims. I shall assess this justificatory practice in section 5.2.1.

Hull's science of science is not a sociobiology of science (cf. chapter 2). Hull argues that gene-based evolution is not crucial as far as the theory that science is a selection process is concerned:

"We are also a social species. Our preference for living in herds is very likely to have some genetic basis...In short, all of the prerequisites necessary for engaging in the process commonly termed 'science' are to some extent programmed into us. However, relative to the generation time of our species, conceptual change has occurred much too rapidly for changes in gene frequencies to have played a significant role. Hence, conceptual change in science may be a selection process, but it cannot be gene-based. Changes in gene-frequencies are liable

[80] D. Hull (1982) p.272-3
[81] D. Hull (1990) p.80-1
[82] ibid. p.80.

to have very little to do with the specific content of particular scientific theories. The mode of transmission in science is not genetic but cultural, most crucially linguistic. The things whose changes in relative frequency constitute conceptual change in science as elsewhere are 'memes', not genes."[83]

The crucial notion Hull uses is that of conceptual fitness, which is concerned with the recognition scientists receive from other scientists, that is, the sort of credit scientists receive when their work, their memes, their ideas, are used by their colleagues. By means of this use of their work, especially explicit citations, scientists increase the number of replicates of their memes in successive generations. Conceptual inclusive fitness is not a biological notion:

"Even if a significant correlation did exist between genetic and conceptual inclusive fitness among scientists, I do not think that very many features of science are going to be explicable on strictly biological grounds. Certainly both the content of particular scientific theories and our understanding of the proper way to go about doing science have changed too rapidly for changes in gene frequencies to have played much of a role either."[84]

Science is not clearly "adaptive" or biologically advantageous (cf. chapter 2). There is no correlation between conceptual fitness and reproduction, since the rise of cognitive status seems to be accompanied by having fewer offspring in "scientific" societies. For EEM the decrease in fertility that is correlated to rise of cognitive status is a theoretical problem (cf. chapters 2 and 3), but it is not for non-genetically based evolutionary theories of science, for EETs. Epistemologies (e.g. EEM) that base their analysis on the phenomenological, perceptual, or more generally psychological "unity of mankind" cannot claim epistemological priority, since, as Hull argues (1988a p. 487), "I see no point in basing an entire branch of philosophy, not to mention all science, on the happenstances of our evolutionary development."[85] Hull (1986) thinks that basing prescriptive philosophy on a contentious postulation is an incorrect strategy. He specifically rejects the claim that there is something that could be called "human nature". The psychological unity of mankind is, in other terms, a

[83] D. Hull (1988b) p.124
[84] D. Hull (1988a) p.283
[85] Even though Hull directs his argument against phenomenological approaches to epistemology, the same argument could be reframed against EEM. As phenomenological and perceptual abilities vary among humans, so do adaptive

residue of the pre-Darwinian essentialism that has tainted our ethical and epistemological analyses (cf. section 2.2).

Hull's science of science is mainly a social epistemology. The emphasis on the cognitive strategies of individuals (typical of EEM and other individualistic epistemologies) is insufficient in order to explain the growth of knowledge. This is because these cognitive strategies are not always individually instantiated. In line with the selectionist perspectives advocated by Campbell (cf. section 3.3), Hull argues that individually instantiated cognitive adaptations are not the only existing "adaptations". Selectionist analyses of sociocultural phenomena embrace a larger adaptationist framework according to which the vicarious social processes of selection create adaptive features that cannot be said to be individualistic, but that more rigorously pertain to the holistic system (i.e. science as an institution). The cognitivism, psychologism and methodological individualism I criticise (cf. sections 2.5, 5.1) completely lack a social dimension, that Hull believes is necessary in order to understand science properly:

"Science is a cognitive process. Scientists must have the cognitive abilities necessary for them to do what they do, but strictly biological considerations do not take us far enough in understanding conceptual change of the sort that goes on in science. Of all the cognitive abilities latent in our species, each person can realize only a small fraction. The notion of cognitive resources is much more relevant to our understanding than innate cognitive abilities. If they are to succeed, scientists must use the cognitive resources which they have available to them, either as individuals or as a group of individuals."[86]

Hull's emphasis on the "pooling" of cognitive resources is entirely analogous to Campbell's emphasis on the economy of cognition (cf. section 3.4). The growth of knowledge cannot be explained without any reference to these social factors. Hull points out that reference to socially transmissible cognitive resources is necessary. Beliefs about the substantive content of science, its methodology and axiology are more likely cultural artefacts than natural (e.g.

---

cognitive capacities. However, in spite of these contingencies, Hull argues, scientists "converge" on similar conceptions of natural phenomena. Ergo, phenomenology and EEM are incapable of explaining such phenomenon.

genetic) products. Hull puts a strong emphasis on the necessity of cooperation between scientists, which cannot be accounted for in biological terms (i.e. reciprocal altruism or kin selection). The nature of scientific cooperation must be explained by referring to the social structure of science, to the nature of science as an institution.

Another general aspect of Hull's EET concerns the nature of the model of scientific evolution he proposes. Hull defends the validity and realisability of a general science of selection. Selection theory can be enlarged as to encompass many natural phenomena because selection is a far-reaching concept. Gene-based evolution is only one among many instances of the process of selection. It is just a contingent fact of history that the first scientific study of selection concerned biological evolution.[87] There is no a priori reason why a theory of social learning cannot be evolutionary. This has a particular relevance since "social learning has been developed to its greatest degree in science".[88] Hull rejects the resistance that exists in some quarters to treating the evolution of behaviour as an instance of the more general evolution of phenotypic traits. It is in principle possible to subject biological evolution and cultural evolution to the same theory, since behaviour is a phenotypic characteristic of organisms and should be explained in the same general adaptive terms. What Hull proposes is a unified framework to treat selective processes. In this respect, Hull makes two basic points. The first is that not only organisms (or even species) are the things that possess adaptations; even social institutions (like the particular social organisation of science we observe) do, as his EET tries to demonstrate. The second point is that the behaviour of scientists is "conceptually" adaptive, in the sense that scientists (largely unconsciously) behave so as to maximise their conceptual fitness. Of course this kind of adaptationism is liable to the usual criticism (Were scientists' behaviours actually selected "for" that particular function? Are not certain behaviours maladaptive?), plus a new kind of criticism, since the concept of adaptation

[86] D. Hull (1988a) p. 284
[87] "The usual objection to consider conceptual evolution from the selectionist perspective is that any departure from the norm (gene-based evolution) is considered a form of deviant selection; but the happenstance that selection in gene-based biological evolution was the first selection process to be studied does not mean that other selection processes do not exist." Hull (1992b)
[88] D. Hull (2001a) p. 21

(perhaps so rigorously circumscribed and defined in biology) is stretched so as to encompass group and population-related cultural properties.

Even though there are no a priori reasons to reject Hull's programme, critics of EET contend that biological and conceptual evolution are profoundly disanalogous (cf. sections 4.5 and 4.6).[89] What Hull (1988a chapter 12) contends is that most of the disanalogies turn out to be vestiges of an oversimplistic and biased view of biological phenomena. The disanalogies are spurious, stemming from a biased interpretation of the scope and explanatory power of evolutionary theory. This is because evolutionary theory is framed in terms of the traditional organisational hierarchy (e.g. genes, organisms and species), which provides an "unnatural categorisation" of evolutionary phenomena, since genes, organisms and species do not refer to natural kinds, given the disparate functional roles they play in selective processes. As a consequence, the traditional categorisation makes behavioural and conceptual evolution look "unstandard".[90] Those who object to extending evolutionary theory to conceptual change usually have an impoverished view of biological evolution:

"According to some theories of cultural transmission, biological evolution is taken to be fundamental, with cultural evolution being only analogous to it [gene-based biological change].... No system has any priority over any other. Biological, social, and conceptual changes are all equally instances of the same sort of process."[91]

Another consequence of taking such hierarchy as fundamental is to have highly contingent and variable selective explanations, hence, no general laws of selection. What Hull proposes is a terminological switch:

"....the regularities that elude characterization in terms of genes, organisms and species can be captured if natural phenomena are subdivided differently: into replicators, interactors and lineages."[92]

---

[89] In particular, it is said that conceptual evolution is Lamarckian, intentional and progressive, while biological evolution lacks these "special" features. cf. ibid. p. 36-40
[90] D. Hull (2001a) pp. 25-26.
[91] D. Hull (1988a) p. 282
[92] D. Hull (2001a) p. 26

This terminological switch will solve the problems resulting from adopting the traditional categorisation, since the new concepts, Hull argues, hopefully pick up "natural kinds". As a consequence, Hull continues, we are in a better position to find proper evolutionary generalisations, or, even, universal laws of selection processes:

"As it turns out, the amount of increased generality needed to accomodate the full range of biological phenomena turns out to be extensive enough to include social and conceptual evolution as well."[93]

By introducing the new terminology, Hull hopes at least to clarify many issues in the philosophy of biology, apart from reaping epistemological benefits. The new terminology, for example, provides an alternative framework to both selfish gene and multi-level selection theory, plus a better solution to the problem of the units of selection (Hull 1988a pp. Chapter 11).

As far as EET is concerned, Hull's primary aim is the application of the new general model of selection to scientific evolution. Replication in science involves different entities: ideas or memes (which, however, should functionally act as biological replicators do). Interaction concerns the activities of scientists and groups of scientists who compete and cooperate among themselves and who test their hypotheses. Interaction in science is a complex process consisting, in a nutshell, of testing, discussion and communication. I shall assess whether Hull's new categorisation benefits his EET. Even though I agree with Hull that none of the disanalogies prevents the possibility of extending selection theory to science, I believe that his new categorisation is epistemologically irrelevant. The issue concerns what a selection theory can say about science. Hull thinks that general laws of selection should be found. Hull's laws are couched in terms of the notion of conceptual inclusive fitness and demic structure of science. Are these laws general enough? We will see in section 4.7.

---

[93] D. Hull (1988a) p. 403

After this very brief illustration of what I deem to be the most important aspects of Hull's general epistemological approach, I move on to the analysis of the mechanism of scientific selection.

4.2 The mechanism of selection

Hull introduces his mechanism of selection in order to account for the general structure of the scientific process (cf. quote at the start of the chapter). The mechanism of scientific selection is based on three elements: curiosity, credit and check. Scientists must be curious in the sense that they must desire to find out about the natural world. This is an axiom in Hull's science of science. Curiosity is essential to science; it is a necessary, but not a sufficient, condition for the growth of scientific knowledge. Other branches of evolutionary epistemology have the role of explaining the nature of this curiosity. Is curiosity an adaptation? Are we, as humans, more curious than amoebae? Than chimps? Adaptive scenarios and putative explanations can be sought. Maybe sociology has some explanatory role to play as well: are scientists more curious than other "social classes", e.g. artists, policemen, politicians, priests? Hull is not concerned about these issues.

Another essential element of Hull's mechanism is credit. Scientists long for credit, they tend to act so as to maximise their conceptual fitness, and in order to do so they need to pass their memes to other generations of scientists. Credit is half the mechanism of scientific selection. The search for credit is considered by Hull as a second axiom, a primitive notion in his theory. Postulating the desire for credit is necessary in order to make sense of scientists' behaviour (Hull 1988a p. 483). Credit is a kind of status, but it is not correlated with biological fitness. The desire for recognition could well be a cultural universal: "Once basic subsistence needs are met, people strive for recognition as for nothing else save sex."[94] Why is this? What is the real sense of this passionate striving? "Money. Respect. Immortality. A way of denying the

[94] ibid. p. 281

randomness that spawned us, and of holding off the fear of death."[95] Of course, a relationship between reproductive and cultural advantages must exist, but, again, Hull is not interested in giving an account of the relationship. This work must be left to other empirical theories. Credit has, of course, a social dimension. In particular, scientists do not primarily seek to gain advantages in terms of social status, but rather seek the more circumscribed credit that comes from their community. Scientists do not usually seek to win prizes, or any other kind of acknowledgement that comes from outside the tight community they form. The basic form of credit in science is use of their ideas, memes, which derives from other scientists citing their papers, their work, from an explicit recognition of their intellectual role.

The third element of the mechanism is check. Credit is half the mechanism, while check is the other half. Checking and testing empirical hypotheses can hardly be seen as a novel element of Hull's theory. But Hull has peculiar ideas about testing and its role in science. The novelty of Hull's views about testing comes from the characterisation of the process: in scientific practice testing is a social effort. Scientists try to refute, disconfirm, falsify other scientists' results, and by means of this activity they enhance their conceptual fitness. Scientists cannot test every result or any previous idea or hypothesis they use in their work; if scientists had to start from scratch every time by testing all the knowledge they incorporate in their research and that they somehow take for granted (being this because of indoctrination, because they deem these results true etc.), science would proceed at an extremely slow pace (perhaps slower than the gradual tempo of biological evolution):

"Scientists learn from their experience, but they do not confront the world of their experience as isolated individuals. If they did, science could never have been developed. Scientists must use each other's ideas, pass them on, and improve on them."[96]

There are two elements of Hull's EET which are particularly innovative. First, Hull thinks of science as a demic process. Scientists tend to organise themselves socially by dividing into

[95] Ian McEwan - "Amsterdam" p.78
[96] D. Hull (1988a) p. 433

research groups or subpopulations. The demic structure is one of those features of the scientific process that have always been treated with suspicion, since normally considered normatively irrelevant. The other novelty of the approach is that Hull analyses science in terms of the individualistic and selfish motivations of individual scientists (and the aims of social entities or scientific interactors, namely, research groups and more inclusive social structures). Of course, while traditional normative views of science see such selfish aims as counterproductive to the achievement of the assumed aim of science (truth), Hull's approach is based on the view that by enhancing conceptual fitness and striving for maximising credit, scientists' behaviour causes the achievement of the institutional goals of science. Hull adds that scientists need not to be conscious to act so as to enhance their conceptual fitness. Nor do they have to explicitly aim at the achievement of truth. Rather, scientists, by means of their multifarious behaviour, cause the process of the growth of knowledge.

The basic process characterising science is, paraphrasing Adam Smith, an "invisible hand" process. Hull is literally proposing invisible hand explanations in terms of unintended consequences: individuals behave to bring about their personal goals and in the process some unintended, serendipitous, emergent effect comes about. The difficulty with these explanations is to specify the nature of the mechanism that "aggregates the dispersed individual actions" (Hull 1997 pp. 118-119). Can the workings of the mechanism be specified? Is there a mechanism at all underlying the process, or is it just due to a series of coincidences? Hull believes that science functions successfully because, most of the time, most scientists behave so as to maximise their conceptual fitness, and, since the credit system in science is structured in a particular way, the behaviour of individual scientists, based on their individual search for credit, fulfils the institutional goals of science in general. What are the institutional goals of science? How is the credit system structured in science? I shall explain this in the following pages.

The mechanism of selection that Hull presents underlines a set of necessary conditions for the growth of scientific knowledge. In order to pre-empt possible criticism, Hull acknowledges that his mechanism has a prominent but not a unique role in explaining the success of science. In fact, Hull acknowledges that certain preliminary conditions are to be satisfied for his "invisible-hand" explanations to be explanatory (scientists "must" be curious, they "must" strive for credit, the causal structure of the world "must" be knowable, society "must" have a certain organisation etc.). In particular, Hull does not say anything innovative about the relationship between theory and evidence. The fact remains that an event (science's success) can have many different causes. The relevant point is rather whether one of these causes is more crucial. Traditional epistemology considers testing as the selective procedure crucial for science to achieve its aim (e.g. truth, accurate prediction etc.). Hull thinks otherwise. His mechanism of selection, he argues, is the most important cause of the kind of success we observe in science. In section 4.6 I shall assess Hull's claim about the causal primacy of his mechanism for science's success.

Hull also adds that the selective mechanism he posits is not universally distributed in science, but that it is more prevalent in some scientific areas than others, and can also be found in certain areas outside science (like academic philosophy of science).[97] However, Hull is committed to the view that science works best at realising its goals when his mechanism is prevalent, and that such a mechanism has a much more important role within science than in other disciplines:

"I intend my account of how science functions to be general: whenever the conditions I specify obtain, the result should be the growth of scientific knowledge.....Nor I am committed to the view that the elements I specify are unique to science. Quite obviously they are not. The sort of mutual citation that goes on in science also characterizes other academic activities, including biblical scholarship......*However, I do claim that the combination of the elements which I specify is unique to science*.....The mechanism I describe does not explain everything about science. Certain preconditions must be met.....If presenting mechanisms to explain why science operates the way that it does is not a proper activity for a philosopher, then so be it. I stand convicted. But if science is to be viewed as an evolutionary process, the specification of a mechanism up to the task is necessary, and this mechanism must be spelled out in sufficient detail."[98] [my italics]

[97] For a criticism of Hull's thesis based on this view, see Kitcher (1988) and Maynard-Smith (1988).
[98] D. Hull (1988a) p. 285

Some critics point out that this view implies a kind of sociological relativism, and that its endorsement does not provide a sufficient basis to solve the problem of demarcation by sociological means. In section 4.8 we shall see whether this is true, and whether Hull tries to solve the demarcation problem at all.

More generally, traditional epistemologists would argue that Hull's science of science is normatively irrelevant. It seems to me that Hull is correct in pointing out that for an evolutionary epistemology it is necessary to describe the workings of the selective process operating in science, analogously to evolutionary biology describing the action of natural selection in all its different forms. Hull also claims that such descriptive and explanatory work has normative value. In section 4.9 I shall return to the issue.

One important point to note is that Hull presents his EET as an alternative to both epistemological externalism and internalism. The alternative view he proposes can be termed "internalised externalism". Hull's conception is thus introduced:

"....one social force that might appear to some to be external to science is internal to science - the sort of cooperative competition that characterizes the social relations among scientists."[99]

Hull argues that the dispute between internalists and externalists has reached a point of no return, since the historical evidence can be interpreted in many incompatible ways:

"Ideally, what one needs is something like a twin study at the level of scientific communities. In the absence of such an unlikely eventuality, it would be nice to find a case in which the predominance of the external factors led in one direction while the predominance of internal factors led in the opposite direction. Which way did scientists opt? Such ideal situations are liable to be extremely rare."[100]

[99] ibid. p. 514
[100] ibid. p. 387

Hull grants that there is some truth in the externalist programme. Externalists argue that the criteria for consensus-generation (often equated to criteria for the validation of knowledge claims) are highly idiosyncratic and "particularistic", in the sense that external causal factors affecting the evolution of knowledge vary in every epistemic situation, as aptly demonstrated by the proliferation of externalist science studies undertaken in the last thirty years. The outcome of this view is that either no "interesting" (i.e. universal) generalisations about the evolution of science can be found, or that not even contingent generalisations can be captured. In a nutshell, what many externalists deny is that science is "a" process (rather than a series of processes). Hull agrees that psychological and social selective "biases" are operative in science, but he resists the consequent inference that no generalisations can be formulated about science. What Hull proposes is to transform externalism. If by externalism is meant the influence of social forces that cannot be accounted for in terms of the workings of the credit or checking system operative in science, that is, influences that are "external" to the sociological structure of science, then Hull thinks that their influence on the process of science has many times been demonstrated to be irrelevant.[101]

Perhaps Hull's views are more consonant to those of internalists. In fact, the trinity of reason, evidence and argument is certainly causally relevant: "....scientists do tend to reach a consensus through the years that is correlated to a convergence of internal but not external factors." (Hull 1988a p. 387) However, Hull argues, those internal factors that led to consensus invariably differ (however slightly) from those that led to the original acceptance of the hypothesis. This means that the criteria of selection that internalists focus upon do vary in time, and also vary from area to area of science. Even though it might be possible that these criteria "......may in the long run be the final arbiters in science often enough for science to fulfil its traditional goals...." (Hull 1988a p. 392), Hull doubts the descriptive and normative completeness of such an approach. This is because science is an historical entity that "evolves", and because empirical evidence is, as a matter of fact, seldom sufficient and always

[101] Cf. Hull (1988a) pp.384-396

liable to be interpreted in multifarious ways in scientific practice. Hull accepts the overwhelming evidence provided by the plethora of science studies suggesting the extent to which scientists act in apparently irrational ways, by interpreting evidence in idiosyncratic manners, by resisting to counterevidence, by conforming uncritically to the received view. Given this situation, internalists invariably will argue that consensus will emerge when empirical evidence is sufficient to adjudicate who is right, or possibly when the confused methodological and metaphysical issues are clarified. This strategy is doubly unsatisfactory for Hull. From a descriptive and explanatory point of view, especially pertaining to the posterior reconstruction of historical episodes in which consensus has been reached, Hull believes that reference to the social structure of science is causally explanatory. The emergence of consensus in the community of systematists as far as the success of cladism is concerned is by Hull explained by reference to his mechanism of selection (cf. section 4.3). From a normative point of view, especially pertaining to those still going-on scientific disputes in which consensus has not yet been reached, Hull believes that EET can provide the basis for extrapolating normative recommendations regarding to best cognitive strategy to achieve consensus formation.

However, Hull knows that the internalist tradition is still miles away from capitulating over the causal relevance of social factors in order to explain such processes. Hard-core internalists still smell the odour of Kuhn's "techniques of persuasion", and they still dread to even consider the descriptive (not to speak of the explanatory and normative) relevance of social factors in situations in which, as Kuhn showed, matters cannot be resolved by "proof". As a result, internalism is bound to leave us with a chronic scepticism about the superiority of science as an epistemic practice, and also with a very poor picture of scientific practice. As Hull puts it:

"In any case, the message [of this book] is that lists of defining properties of science are of secondary importance to the continued action of the mechanism that I have set out. If there is an essence to science, this mechanism is at its core."[102]

---

[102] D. Hull (1988a) p. 392

Hull agrees with internalists but disagree with externalists that science is a practice and that generalisations can be captured about the scientific process. Furthermore, from such generalisations normative recommendations can be derived to make science more effective in achieving its aims (e.g. consensus formation and objectivity). Second, Hull disagrees with both externalists and internalists because they refer to idiosyncratic causal factors to explain science's success. In particular, internalists refer to criteria of selection that evolve (apart from differing from discipline to discipline) and change, and upon which we cannot generalise. The only factors upon which we can generalise are those concerned with the social structure of science. In sections 4.6 and 4.8 I shall focus on the epistemological relevance of Hull's social epistemology. Hull's EET would be vindicated if he managed to pinpoint the social factors that are effectively causally relevant in consensus emergence and knowledge validation. I think Hull manages to achieve this aim only partially. As a consequence, I think that Hull's EET is not generally a better alternative to internalist and externalist epistemologies, even though it is part of a more complete epistemology of science. Still, Hull contends that his "internalised externalism" is alternative to both approaches. This is because the defenders of these approaches do not realise that some social features of the process of science are adaptive and adapted, that reference to such features cannot be eschewed to explain scientific change, and that proper generalisations about the scientific process must refer to them. These features pertain to the social structure of science. I shall now move to assess the adaptive nature of the scientific process, which provides the core of Hull's EET.

4.3 The nature of the scientific process: competition and cooperation

As Hull puts it, science is the interplay of competitive and cooperative scientists' behaviours. A kind of Darwinian struggle for conceptual survival with some episodes of altruism constitute the path of science. Both competition and cooperation play an essential part in the process, hence an analysis of science that eschews reference to either of these processes is incomplete. The major thesis of Hull's EET is that the competitive-cooperative organisational structure of science is "adaptive" and "adapted". Science is competitive because scientists are motivated by the fulfilment of apparently selfish needs: curiosity, credit and, negatively, by checking. The peculiarity of the process is that the (usually selfish and) multifarious motivations of scientists tend to have as the outcome the realisation of the institutional goals of science. Science is cooperative because scientists are rewarded for cooperating. In this sense the credit system of science is so structured that scientists can be seen as forced to cooperate in order to receive credit for their work. As a consequence, such cooperative behaviour is essential for the growth of knowledge.

That science is a highly competitive activity is clear to everyone familiar with the disputes, debates and diatribes punctuating its history. Science is competitive because scientists need to behave in particular ways in order to enhance their conceptual fitness, in order to spread their memes and receive credit from their competitors. For Hull, building alliances with other scientists (e.g. by forming a research group) can positively enhance their fitness. Inspiring allegiance from one's peers is another strategy for fitness-enhancement. However, aggressiveness in proposing and defending one's ideas, and other supposedly negative behavioural strategies, can also have positive effects. In practice, Hull argues, scientists resort to all manner of "pseudo-criminal" behaviour in order to enhance their fitness: they struggle for the editorial control of journals, they tend to obstruct publication of unsympathetic research, they show various sorts of bias, and a great preoccupation with power and control within the profession and the academia.[103] Scientists are generally not motivated by the ideal of seeking knowledge for the sake of truth but, rather, by various personal commitments, self-interests and political motivations. The desire for credit accounts for the polemics in science, for the distortions and inattentions scientists give to other scientists' ideas, for episodes of theft and fabrication, in brief for all those putative disreputable features of scientific practice. Scientists promote ideas aggressively in the various academic and non-academic arenas

[103] As Allen (1991) puts it: "Nice guys don't win Nobel Prizes".

(specialist journals, professional meetings, review panels, university departments etc.); they compete for their ideas to be replicated more than those of their competitors, in order to leave more copies of their memes in the next generation.

The pursuit of credit is constrained by the possibility of checking. With the possibility of apportioning credit (and the subsequent enhancement of their conceptual fitness), comes the possibility of apportioning blame (and the subsequent reduction of conceptual fitness). Scientists are motivated "negatively" by checking in the sense that by showing that, for instance, a particular experiment, when replicated, does not yield the expected outcome, they will blame other scientists, enhancing their fitness and decreasing the fitness of their competitors. Pons and Fleishman experienced the efficacy of such a blaming mechanism when their results on cold fusion could not be replicated. In Hull's (1988a p. 353) colourful language, the "get that son of a bitch effect" is a major motivation of scientists.

Scientists are not necessarily aware of why they are behaving in particular ways, exactly as organisms are unaware of struggling to maximise their genetic inclusive fitness. Not only are scientists not conscious of behaving so as to maximise their conceptual fitness, but they are even less aware that their selfish behaviour brings about the institutional goals of science in general. Furthermore, Hull (1988a pp. 319-320) suspects that scientists actually had better not be aware of their struggle for fitness, since, if they knew, this reflexivity or feedback process could affect negatively the nature and success of the overall process of scientific evolution.

An important characteristic of Hull's account of competition in science is that it is ramified at different levels of the social structure of science: it is not only individual scientists that compete, but also groups and other more inclusive social entities. Intra-group and inter-group competition are both paradigmatic cases. Competition for funding arrives to encompass different departments, scientific disciplines, entire universities and organisations.

One of the basic peculiarities of the scientific process is that scientists are "forced" to cooperate with their closest competitors because they must use each other's work and cognitive resources. The cumulativity of science partly depends on this kind of egotistical and forced cooperation. That scientists are able to build on the work of their predecessors is a necessary condition for the growth of knowledge and the cumulativity of science.[104] Hull claims that:

".....science can be viewed as a self-policing system of mutual exploitation - or cooperation if one prefers."[105]

The system of mutual exploitation is based, yet again, on the quest for credit. There is a need to cooperate with one's competitors because of the structure of the credit system. In order to receive credit (and increasing their conceptual fitness) scientists need to cooperate by giving credit where credit is due (for instance in using the work of others in order to develop their hypotheses). The process results in a conflict of interest: citation automatically enhances others' conceptual fitness and detracts from originality, but not apportioning credit where credit is due can have major deleterious consequences. Scientists share credit and recognise those who helped them because doing otherwise would lead, in extreme conditions, to isolation and exclusion from the scientific community (and not being part of the community means having a very low conceptual fitness). As a consequence, Hull remarks that the upshot is that scientists will cooperate only when it will be useful to them.

Particular emphasis is put by Hull on another form of forced cooperation. Crucially, scientists form alliances, however opportunistically, in order to compete better with opponents. The existence and functional role of such professional relationships is essential, Hull argues, in order to understand the presence, so typical of science, of matters of inter-group hostility and intra-group loyalty. In science, Hull continues, the tendency is for scientists to create groups

[104] By cumulativity Hull means the process by which, given a problem, successive theories tend to solve it better and better. Hull has doubts about the extent to which science is cumulative. But the point remains that cooperation is necessary for science to be cumulative in the sense he defines.
[105] D. Hull (1996) p.72

or conceptual demes. The demic structure of science has a functional role, and provides benefits to the various interactors operating in science (individual scientists and research groups). The group plays as important a role (that of, by using Hull's terminology, an "interactor") in scientific selection as the individual. It is as important as the isolated individual because "sharing a common fate" by collaborating will enhance the conceptual fitness of all the scientists belonging to the group and also the credibility of their project. The benefits to individual scientists given by the demic structure of science are plentiful. Forming alliances has two very important functions. First, the members of the alliance can share conceptual resources, given that many problems require an array of cognitive resources that no single scientist can possess. Second, members can share more easily professional resources of different kinds, like the control of academic societies, the control of the editorial policies of professional journals and the influence on funding agencies. The members of the group will use and cite each other's work more frequently. Positive citation promotes social cohesion and possibility of replication of scientists' memes, with consequential increase in visibility of their papers, possibility of gaining more graduates and more financial resources (all measures of fitness-enhancement). The demic structure of science also facilitates testing. Members of the alliance will check the results of colleagues, reducing the risks involved with the publication of unprepared research. Another basic advantage of the demic structure is that the amount of intra-group cognitive variation is increased, with the consequence that science (that, by assumption, is a selection process) can reach its aims more quickly.[106] In fact, remember that ample variation is a necessary condition for selection processes to create adaptations (cf. section 4.5).

Of course, such a social organisation has potential disadvantages. One danger of this organisational feature of science is that inter-group warfare can become too intense, with the consequence that factionalism can affect and influence various mechanisms of selection. As always with selection processes, Hull argues that an equilibrium-state must be reached between potential benefit and disadvantages. For instance, the demic structure of science will

[106] Perhaps Hull would agree that in order to further increase conceptual intra-group cognitive variation we should

have a functional role in furthering the institutional aims of science only if a stable relationship can be reached between factionalism and inter-group cooperation. Hull claims that there are no general recipes for showing how such equilibrium-states can be reached, given the role contingency plays in science. In this respect, Hull argues (1988a p 507-8), chance is very important in science (as drift is important in biological evolution).

A basic element in Hull's analysis of science concerns the more general sociological structure of science as a process: science is a self-policing system. Hull's claim is that science is the most successful epistemic human practice because it is designed as to make the interests of individual scientists largely coincide (despite the multifarious arrays of their personal aims and motivations) with the interests of science (its institutional goals). Science is structured in such a way that the vast array of adaptive social mechanisms (credit, competitive testing and the demic structure of science) will act in order to curb individual scientists' and group's biases. In this sense the role of group cooperation is also crucial. Members of the group, for instance, diminish the possibility of fabrication and fraud. The most important point is that self-policing systems function properly only when individual and group interest coincide. Scientists, differently from doctors and politicians, police themselves well because it is in their own interest to do so. Hull's point is that it is not necessary to consider scientists as driven by a univocal desire for truth, to unrealistically conceive of scientists as "saints", or, as Medawer puts it "amiable creatures".[107] As Hull reminds us (1997 p. 123, following Hume?), "the best security for the fidelity of mankind, is to make interest coincide with duty". This is exactly what happens in science, since scientists' competition for credit is constrained by mechanisms that make the game as fair as possible. More generally, Hull's claim is that the structure of science is adaptive (i.e. it is historically a product of selection processes) and largely adapted (i.e. it is designed to maximise the achievement of its institutional aims, namely the growth of knowledge). The claim about the adaptedness of the system is to be

advise scientists to have a philosopher for research group!

[107] Medawer also claims that: "There must be very few wicked scientists. There are, however, plenty of wicked philosophers." cf. Hull 1997 p.122

understood properly. The self-policing of science functions properly in certain conditions, but forms of maladaptive behaviour are possible, especially if, to draw an evolutionary analogy, the "environment of ancestral adaptation" (that to which science is adapted) changes. Hull is conscious of the potential role that external and interfering influences have in affecting the adaptive nature of the scientific process:

"As scientists are increasingly able to make money off their discoveries, the same sort of financial impropriety that characterizes all other professions will increasingly characterize science. Whenever scientists serve two masters, compromises will be made, whether these masters are government, industry or mammon."[108]

Science policy makers have the duty to come up with policy solutions as to re-establish the effectiveness of the self-policing system.[109] But the task for science policy makers is difficult, because it is too easy to come up with facile recipes without taking into account that science is a creative process:

"What are the optimal levels of selection in science if what we want is a maximum production of high-quality work? We do not know. One thing that we do know about selection processes is that they are inherently inefficient. Because contingency plays such a large role both in the generation of novelty and its selection, large numbers are required. Small numbers increase the likelihood that new species will evolve, but they increase the likelihood even more that these incipient species will go extinct. As a result, *attempts to increase the efficiency of selection processes are likely to be quite dangerous*. Calls to make science more efficient might sound plausible enough, while comparable calls to make the production of great works of art more efficient might not sound quite so sensible. Science is inherently a creative process, and creative processes are inherently inefficient. One explanation for this inefficiency is the large role that selection plays in any creative process."[110] (my italics)

Another reason why Hull claims that it is difficult to devise an alternative organisational structure for science is that he believes that the interplay of cooperation and competition (properly constrained by the specifically designed reward and punishing mechanisms) works generally fairly and sufficiently well. For this reason, critics argue that Hull is defending the

[108] D. Hull (1996) p.72
[109] This is one of the potential ways in which EET has a normative role to play.
[110] D. Hull (1992b) p. 75-6

status quo, assuming that science is already suited to achieve its institutional aims, and by accepting the deficiencies and distortions of the scientific process as an inevitable but still "natural" outcome (cf. section 4.8). In particular, Hull realises that too much competition can be deleterious and maladaptive, or that too little cooperation can have similarly negative effects. But altering the already adapted structure of the scientific process might have even extreme consequences:

"If the self-interest of scientists is redirected, then the basic structure of the scientific enterprise is likely to be modified and the results altered accordingly. We cannot take the structure of science for granted. For a long time it was largely absent from human affairs. It came into being; it could just as readily cease to exist."[111]

Hull seems to think that most of the time the self-policing organisation of science will work naturally as to reach a new level of adaptedness, a new equilibrium state.[112] Thus, Hull's major claim is based on the contention that in the long run the selective process of science will lead to the selection of the most-valuable scientific ideas.

I have so far given a general outlook of Hull's thesis. In the remainder of the section I shall provide three examples in order to illustrate Hull's epistemological work.

The first illustration pertains to the adaptive explanation Hull offers to show why lying is rarer than stealing in science. Contrary to what happens in "normal life", where theft is generally considered to be a more damaging vice than lying (and lying is certainly more prominent), theft is more common than complete lying in science. One clarification needed pertains to the varieties of lying involved in scientific practice. In this respect, Hull argues that there is a fine line between accomodating evidence, proper lying and complete fabrication. Hull

[111] D. Hull (1988a) p. 392

[112] The concept of equilibrium state can be characterised in many ways. It generally refers to a stable relationship between different quantities. The most general sense concerns the most epistemically conducive levels of intellectual variation and retention. A less general sense concerns the relationship between conservative and innovative or deviant behaviour as far as selection is concerned. An even less general sense concerns the optimal levels of cooperative and competitive scientific behaviour. Hull also refers to the most effective level between intra-group cooperation and factionalism.

(1988a pp. 315-319) thinks that accomodation of data sets (what he calls "finagling" or "sanitizing" of empirical evidence) is, in practice, a necessary element in presenting and defending one's hypotheses, a necessary element of science. Mendel, Newton, Galileo, and many others among the greatest scientists have been blamed for the extent to which their observational reports fitted their theories. On inspection, such fit could be seen as accomodation. Among the possible kinds of lying, failure to give credit is less common than finagling, but far from being uncommon. Scientists tend to accentuate the originality of their proposed ideas, failing to give credit to some peer or precursor (and especially to opponents of their research programmes). Concerning stealing, Hull claims that it is possibly endemic in science: which graduate student is ever going to object to her "master" for stealing (perhaps with care) her ideas?

According to the official ideology of science, scientists "should" not tell lies, they "should" not fabricate empirical results, they "should" not steal ideas. These are norms of science, violated in practice. A descriptive analysis of the practice of science shows, however, that fabrication is rarer than stealing. In order to explain this fact Hull claims that in science the social system works (and should work) as to minimise fabrication. The social mechanisms designed to minimise the amount of fabrication will lead to a greater punishment than in cases of theft. This is because stealing hurts only one scientific interactor, while fabrication hurts the whole community of scientists. Fabricated research hurts not only one's competitors but also one's allies. Outright fabrication harms every scientist who relies on that research. Hurting one's opponents is better than hurting one's allies. For this reason, mechanisms for the detection of any form of lying must apportion different kinds of blame depending on the extent and gravity of the lying process involved. The scientist who adjusts excessively his theory to the data, or the scientist who invents a set of confirmatory data just in order to provide confirmation for his hypothesis, affects much more the cumulativity of the selective process than the scientist who robs ideas from, for example, his graduate students and his less important peers. Hull's sociological explanation refers to the existence of evolved selective mechanisms that act as to maximise the possibility of detecting error, and as to minimise the possibility of error ramifying. Hull's explanation that lying is rarer than stealing starts from the assumption that the social organisation of science evolves in order to maximise the effectiveness of the cumulative process of knowledge acquisition. The adaptive system of science is so structured as to punish more severely forms of behaviour that affect negatively the cumulativity of the scientific process. The cogency of this generalisation shall be assessed in section 4.7.

In order for the interplay of the action of the credit and testing mechanisms to function properly, and contribute efficiently to the growth of knowledge, inter-subjectivity of research is needed. Science has evolved mechanisms specifically designed to force scientists to publish their results and to improve the communal access to knowledge. One such mechanism is the following. The antithesis to communal access to knowledge is secrecy. It is therefore not very surprising, from Hull's perspective, that secrecy in science is generally constrained, since it is "maladaptive". A basic property of science as it is practised today is the convention at the heart of the credit system in science, where the first publisher of some research takes all the credit. This feature of science seems to collude with the idealistic view according to which scientists seek knowledge for its own sake. If only scientists were capable to refrain from their urge of enhancing their status and were more prone to act as to maximise the general good, perhaps such a convention could be given up. But, as Hull stresses, this is not so. In fact, such a convention is an adaptive feature of science that influences the way in which scientific change occurs. This convention, Hull explains, has a functional role: it "forces" scientists to publish.

Hull's sociological explanation is based on the tenet that the credit system in science is organised so as to minimise such form of maladaptive behaviour. Since scientists would like to keep secret their innovative results as much as possible (in order to derive their immediate consequences and remain a step ahead in the competitive struggle they face), a mechanism has evolved with the effect that if they want credit they must publish. Hull's EET not only explains why such a convention exists and why science has this functional property, but it also predicts that to give up the convention would have a deleterious effect on scientific evolution (i.e. the elimination of such a convention would hinder the cumulative nature of the scientific process). In order to disconfirm Hull's EET we thus need to show that the credit system in science can be organised more fruitfully by giving up the "winners-take-it-all-convention". The alternative hypothesis claims that science could be organised even more efficiently by giving up this social feature (e.g. by sharing credit between the members of the scientific community). However, it seems that the available historical evidence confirms Hull's hypothesis about the adaptive character of the credit system over the alternative. In fact, in 1666 the French Academy was formed by endorsing the convention of communal credit (i.e. credit for publication of papers was given directly to the society). In the space of 33 years, the Constitution of the Society was changed, embracing the individual credit convention (Hull 1988a pp. 322-3). This is because the system did not work: scientists complained about the lack of reward, and secrecy increased.

As far as the adaptedness of this feature of science is concerned, as usual matters are about relative functionality. In practice, respect of this convention results in the sometimes deleterious outcome that is the rush to publish unprepared papers, and in the increased incidence of priority disputes. All these features can be seen as maladaptive, as is the secrecy of knowledge. What Hull (1996) contends is, in this respect, that the rush to publish, when properly constrained, increases the pace of science. If priority disputes and publication of bad research is going to increase substantially, then some policy measures will be needed (e.g. by acting on the constraints on publication etc.). But the point remains that, as science has been practised until today, the convention has had a positive effect on the growth of knowledge, and that communal access to knowledge is much more beneficial both to the eventuality and risk of increased priority disputes, and to the publication of unready research.

One of the innovations in Hull's EET is the reference to the demic structure of science. In order to illustrate Hull's argument about the functional and adaptive role of the demic structure of science, let us consider his main test case. Hull explains the success of cladism over other taxonomies as being significantly caused by the fact that cladists were, as a group, more cohesive and tight. Two kinds of factors influenced this success. The first pertain to the personal capability of individual cladists. For example, Farris' and Nelson's aggressiveness in proposing and defending their ideas, or Farris' politicking and ability to understand that the struggle for the editorial control of journals, their refereeing process, and academic societies, was crucial. In this category we can also put Platnick's and Nelson's ability to choose the best "patron saints" in order to substantiate their claims. For instance, they choose to form an "alliance" with one of the most important and most respected "authorities" in the philosophy of science (i.e. Popper). The second kinds of factors that Hull individuates are more importantly group properties. Cladism won the argument with pheneticism and evolutionary taxonomy because cladists were much better at networking and at building professional allegiances and alliances. For example, cladists were more effective in defending the nature of their endeavour by means of all sorts of intra-group cooperative and inter-group competitive behaviour: speaking with one voice when one of the tenets of cladism was attacked and criticised from the outside (i.e. by their rivals), and also publicly defending members from such criticism. In addition, they maintained public agreement on fundamentals despite internal conceptual heterogeneity, while removing anti-conformist and "heretical" members from certain crucial academic or editorial positions. Members' attempts to tie cladistics with potentially highly criticisable theoretical engagements (e.g. Nelson's association of cladism with Croizat's biogeography, or Wiley and Brooks' connection with thermodynamics) were also resisted. The essence of their success, as Hull sees it, is that cladists promoted conceptual orthodoxy among the members of their group much more strictly than pheneticists or evolutionary taxonomists did. The defence of this orthodoxy was instrumental in keeping them focused on the main aim to be achieved, that is, securing their primary position as the best theory in systematics (Hull 1988a chapter 7). Hull explains the success of cladism by referring crucially to the beneficial effects that intra-group cooperation had on the achievement of their communal aim. Professional alliances between cladists were much stronger than between, for example, pheneticists. As a consequence, individual cladists reaped more benefit from such networking.

An important point that needs to be noted, Hull argues, is that the amount of conceptual homogeneity among cladists was not more than that among other competing groups. In this respect, Hull stresses that intra-group conceptual homogeneity does not seem to be an essential element in the explanation of the success of cladism. What is rather essential is the intra-group loyalty despite cognitive heterogeneity. For instance, the maintenance of purely professional relationships can account for why Farris chose to stick with Nelson and Platnick despite not agreeing with them on substantial theoretical issues. As Hull says, Farris' choice can only be understood in terms of the benefits his position gained by respecting the alliance with Nelson et al. Cladists certainly did not behave as scientists should. The analysis of the refereeing process Hull provides shows an increase in factionalism while cladists had editorial control of the academic journal "Systematic Zoology". However, such factionalism was not too intense. As a consequence, Hull argues, an equilibrium between inter-group warfare and intra-group loyalty was reached in this case, and science was allowed to achieve its institutional aims. Such an achievement can be properly understood only by reference to the demic structure of science. As I have said before, Hull claims that one expects the effect of the demic structure on rates of conceptual change to be analogous to what happens in biological evolution: as in biological evolution speciation occurs more rapidly when species are subdivided in partially isolated demes, so in conceptual evolution small research groups who pool cognitive resources tend to have a beneficial effect, increasing the rate of change (Hull 1992b pp. 76-8). Hull's test case provides evidence of the influence of the demic structure of science on scientific evolution. From the evolutionary taxonomy's dominion over the science of systematics characteristic of the 1950s, we pass to the dominion of cladism in the 1980s. To be noted is also the fact that pheneticism quickly rises to prominence and then almost disappears in a very short time. Besides, the analogy with rates of speciation borrowed from biology seems, Hull contends, further confirmed by the fact that cladism split in two groups, eventually leading to the emergence of two distinct research programmes (i.e. pattern and phylogenetic cladism).

Another important aspect of Hull's explanation is that it is a group-selection hypothesis. As in biological evolution, group selection is the selection of populations because of some properties that the populations exhibit, so in science, group selection is caused by some property that the group has as a group, independently from the properties (and fitnesses) of the individuals constituting it. In Hull's EET fitness is also a property of groups and larger social entities. Whether group properties are reducible to those of constituent individuals is an important epistemological issue. If groups and their properties are causally relevant for the explanation of cases of epistemic success, then methodologically individualistic epistemologies are wrong (cf. section 5.1). Hull believes that proper (i.e. irreducible) group selection happens in science (as it happens in biological evolution – Sober & Wilson 1998 esp. chapters 1 & 2). Cultural group selection is a reality in Hull's framework.

Given the importance that Hull's test case plays for his science of science, I shall dedicate much ink to its assessment. In section 4.6 I shall defend group selection hypotheses, while at the same time contending that they only provide additional explanatory reasons to understand the history of science. In section 4.8 I shall argue that Hull's EET seems to me much stronger in offering generalisations concerning the scientific process as a whole rather than generalisations concerning particular historical episodes.

In Hull's EET research groups play a crucial role as scientific interactors. In the following section I shall focus on Hull's contention that rationality is not significantly a property of individuals, but of groups.

4.4 The social and evolutionary character of rationality

One of the major features of Hull's epistemology is the emphasis on the social character of the process of knowledge acquisition. As a consequence, Hull also claims that the evolutionary perspective has some consequences as far as the notions of scientific objectivity and rationality are concerned:

"Science is inherently and necessarily a community affair. Certainly isolated hermits can learn about the world, but if science had been constituted in its early years by such hermits it never would have gotten off the ground. In order for science to be cumulative (to the extent that it is), transmission is required. Similarly, the sort of objectivity and rationality that gives science the peculiar features it has are characteristics not of isolated individuals but of individuals cooperating and competing in peculiarly organized groups."[113]

Hull is arguing that rationality and objectivity are (also) properties of populations. This point was already stressed by Kuhn, when he showed that the emergence of consensus (by means of paradigm switches) could not be accounted for in terms of traditional notions of rationality. As Hull puts it:

".....all Kuhn has claimed is that simplistic analyses of rationality cannot explain such transitions. A community-based notion of rationality is more appropriate."[114]

Kuhn's views about the nature of rationality have been criticised because he apparently transformed the process of transition from one paradigm to another in an arational conversion process. In my opinion, this criticism is valid only if it is presupposed that the standards of rationality are universal. What Kuhn had to show was that a different model of rationality has to be endorsed. He therefore proposed to consider the paradigm-dependent community of scientists as the locus of rationality. In Kuhn's view, judgements about the rationality of particular individual behaviours should be relativised to standards of epistemic evaluation that

[113] D. Hull (2001a) p. 46
[114] ibid. p. 45

are paradigm-dependent. In the end, what is rational is only the decision taken by the community of scientists, which ultimately constitutes the best criterion of scientific objectivity. Consensus emergence, Kuhn seemed to claim, is the epitome of rational action. Hull considers this alternative model of rationality unsatisfactory. This is because Kuhn thought of paradigms as homogeneous and monolithic conceptual systems, where members must share specific and univocal methodological and theoretical commitments. They must therefore share a conception of what constitutes rational action. In fact, in Kuhn's view the switch process becomes a matter of giving up the notion of "rationality" as prescribed by belonging to a particular community, and to endorse a different notion that is still paradigm-dependent. Kuhn criticised the traditional universalistic and individualistic view of rationality, but his view on the matter retained, according to Hull, an essentialistic or absolutistic element that must be given up. In Hull's opinion, Kuhn did not understand that even his community-relative (i.e. paradigm-dependent) notion of rationality is still too simplistic. In fact, according to Hull, Kuhn's "paradigms" are, as biological species, individuals. Such historical entities constantly change their nature. In this sense, it is not a problem if communities are constituted of individuals who disagree with each other in many important respects, and whose conception of what constitutes "rational" action diverge. After all, this heterogeneity is, from the evolutionary perspective, a strength rather than a weakness. Hull's notion of group or social rationality presupposes the existence of divergent opinion, which does not prevent the emergence of cooperative behaviour, which Hull considers to be essential for knowledge growth. As Hull's main test case illustrates, genuine episodes of scientific progress happen even though there exist intra-group theoretical heterogeneity, and possibly intra-group disagreements about individualistic or paradigm-relative notions of rationality. In Hull's test case consensus was reached even though Farris, Nelson and Platnick did not endorse the same theoretical, methodological and metaphysical tenets.

I believe that Hull's conception of social rationality is still vitiated, like Kuhn's, by an important problem, since, for instance, the emergence of consensus in the community of systematists cannot be equated with the claim that the success of cladism is a rational process. In this sense, Hull's and Kuhn's notions of social rationality and objectivity are not very different, since the bedrock of both approaches is that social consensus can be seen as evidence of rational action. By considering rationality a social rather than an individualistic affair we reap some epistemological fruit only if we agree that the judgement of the community of scientists is considered the best criterion of objectivity. But this thesis cannot be justified simply by arguing that it is better to trust the many than the few. For this reason, both Kuhn's and Hull's notions of social rationality can be criticised because they do not secure a link between the emergence of consensus in the community and the rationality of such a process. The problem concerns the quality of community consensus, which is far from being *per se* trustworthy. As a result, Hull's views are as relativistic as Kuhn's.

Both Kuhn and Hull want to deflate the notion of individualistic rationality, and substitute it with a social analogue. But the problem faced by Hull remains very similar to the problem faced by Kuhn: why does the ramification of cognitive variation within the community (e.g. between individuals, within research groups, and also within other larger social interactors) make the scientific process (e.g. paradigm shift, consensus emergence etc.) more rational? Why do we need a notion of social rationality at all?
I believe that the solution to this problem consists in enlarging the perspective from which to evaluate claims about rational action. It seems to me that the best move open to EET is to argue that rationality is a property pertaining to the scientific process as a whole, rather than to individuals, specific research groups, or communities of scientists. This view has been advocated by Campbell, who thus articulates it. First, he claims that selection theory is involved in disputations about the rationality of belief-change only insofar as the process of variation-selection is regarded as rational (Campbell 1981 p. 512).[115] The rationality of the scientific process as an instance of a selection process is still based on the "dogma" of

[115] This strategy should not be interpreted as meaning that questions concerning individual rationality are epistemological irrelevant, but only that EET as a social and evolutionary epistemology is largely irrelevant to the construction of these models. This view is compatible with the epistemological pluralism advocated in ch.1

Campbell's approach, namely that any inductive achievement is an instance of a selection process, and the further assumption that if such a process is aim-conducive then it is "by definition" rational (notion of instrumental rationality - cf. section 1.2). Second, he embraces a much more holistic perspective from which to judge whether individual or group selective behaviour is rational. The model of rationality Campbell proposes might be termed "evolutionary" rather than social. More specifically, Campbell argues that certain patterns of scientific behaviour can be properly understood only from an evolutionary perspective that takes the whole scientific process into account. For instance, resistance to counterevidence, insensitivity to counter-instances, and other endemic features of the practice of science that according to traditional epistemology can only be seen as biases and instances of irrational behaviour, can be interpreted as essential features of the scientific process. These biases are sometimes equated with "miscarriages of the scientific method". Kuhn showed that they are typical of science, and that a complete descriptive and explanatory theory about the scientific process cannot elude reference to this kind of behaviour.[116] What Campbell adds is a normative thesis. From an evolutionary perspective, he argues, science is and should be organised so as to minimise the effect of such biases; however, such biases cannot, and, more importantly, should not be eliminated completely because they are an inevitable and essential outcome of the social organisation of science. In this sense, Campbell refers to the adaptive system of science as a "vehicle" that carries knowledge. The structural properties of such a vehicle are partly adapted to the achievement of the institutional aims of science. Therefore, since the social structure of science is causally efficient in this endeavour, it needs to be protected. As a result, many kinds of deviant behaviour can be rationalised if they are seen as instances of "vehicle maintenance" behaviour, that is, as attempts to protect the adaptedness of the scientific institution. Campbell provides a vast array of analogies by means of which he is able to substantiate this thesis, especially those about the doubt-trust ratio (Campbell 1977 pp.477-482). In a nutshell, "deviant" behaviour can be "rationalised" from the evolutionary

[116] There seems to be empirical evidence that attitudes of oversimplification and conservative behaviour are endemic even in children's knowledge acquisition - cf. Gopnik (1996).

perspective Campbell endorses, even though such rationalisation can only be achieved retrospectively.

One by-effect of this evolutionary perspective is that its deflated notion of rationality seems to solve Kuhn's and Hull's problem of linking emergence of consensus with objectivity. In fact, much of Campbell's analysis is directed to show how the emergence of consensus, opportunely constrained by various vicarious selectors, will possibly determine an hopeful augmentation of fit between world and knowledge. In particular, what he contends, and I am not sure Hull agrees with this (cf. section 4.8), is that the sociological "norms" of science are designed to channel consensus so as to maximise the opportunity of "reality" to optimally influence the nature of knowledge. If this is so, EET has a normative role in providing recommendations acting as to protect the adapted features of the scientific process.

By taking conceptual fitness and the demic structure of science as basic properties of the scientific process, what do we understand about science, and how is this understanding different from that provided by traditional epistemologies? In the following sections I shall assess the epistemic implications of Hull's EET, in particular its adaptationism and selectionism.

4.5 The nature of intellectual variation

The explanatory power of selective explanations depends on three conditions: richness of variation, non-directedness of variation, and non-purposive nature of the selection processes involved (Amundson 1989). The first condition is satisfied in scientific evolution: the variety of memes, ideas, hypotheses, theories and research programmes that compete in the struggle for survival and replication is huge. Furthermore, EET supports the view that cognitive variation affects the evolution of science much more than allowed by alternative epistemologies, and, more importantly, that this "should" be so, given that a precondition for science to be successful (in the sense of increasing the fit between our picture of the world and the world itself) is the abundance of variants for selection to choose from.

We now need to assess whether the second condition is satisfied. In section 3.4 I proposed Campbell's arguments to show that the condition is satisfied. EET is committed to the validity of the model proposed by Campbell, which provides the only available explanation of the workings of creative processes by pointing out that scientists' psychology is just a mixture of chance and preliminary selection, and which rejects what I termed the Lamarckian challenge. In this section I shall assess Hull's answer to the challenge.

The Lamarckian challenge is a double-edged sword. On the one hand, it negatively affects the possibility of giving a Darwinian account of knowledge growth; on the other Lamarckism (i.e. the thesis of the directedness of variation) seems a necessary condition for "rational" selection, and therefore a necessary element in explaining scientific progress. In both cases it prejudices EET's epistemic chances, by negating either its validity as an "evolutionary" (in the sense of Darwinian) epistemology, or its capacity of explaining scientific progress. How to reconcile these two incompatible effects? What Hull tries to show is that, first, it is simply confusing to call conceptual evolution Lamarckian, and, second, that scientific growth can be explained by means of the variation-selection model.

The Lamarckian challenge stems from two main sources. First, according to Lamarckism there exists a mechanism for the inheritance of acquired characters. In this respect, Hull contends, conceptual evolution is not in any interesting sense Lamarckian. After all, nobody claims that the transmission of information about the environment through social learning depends on a mechanism of genetic inheritance. In fact, parents transmit ideas non-genetically to progeny (or horizontally), so that ideas are not genetically assimilated by the offspring. Scientists face problems and try to solve them as organisms face environmental problems and try to solve them. Scientists try to find solutions to these problems as organisms modify their phenotypes, behaviour and "habits" (the term used by Lamarck), to respond to environmental challenges. In both cases, a process of ontogenetic adaptation is triggered by the environment. One difference between biology and science is that the process of ontogenetic response has no phylogenetic consequences in biology. For this reason today we say that Lamarck was wrong about the existence of a mechanism of the inheritance of acquired traits (this mechanism does not work because somatic cells do not transmit information to germ cells - this is the central dogma of molecular biology). In science, the process has doubtlessly phylogenetic consequences, but only because the mechanism of inheritance has a different nature: scientists can pass on their memes by means of books etc. to later generations. If the existence of a non-genetic mechanism of inheritance is sufficient to call the scientific process Lamarckian, then so be it. Science might be a Lamarckian process, but, Hull contends, only in a purely metaphorical sense. This sense does not, by itself, affect the possibility of treating EET as a Darwinian theory, in the same way in which biological processes that are governed by non-genetic inheritance of acquired traits (and that for this reason could be called "Lamarckian") are explained in Darwinian terms. The evolution of behaviour is just a case in hand: if an organism finds an adaptive response to a new environmental challenge, and if such knowledge is passed on to other organisms, either because other organisms are able to imitate, or because the original organism was able to transmit such knowledge socially, then the trait can evolve by natural selection. Scientific evolution is Lamarckian just as the evolution of behaviour is Lamarckian. That is, the mere fact that non-genetic mechanisms for the inheritance of acquired traits are operative is not a sufficient reason to call such processes Lamarckian. This is a purely terminological choice:

"In the absence of anything like the [genetic] inheritance of acquired characters, I think that characterizing conceptual change as 'Lamarckian' leads to nothing but confusion."[117]

The second, and relevant, sense in which conceptual evolution can be said to be Lamarckian concerns not the mechanism of inheritance of acquired traits, but the directionality of intellectual variation. The analogy between scientific and organic evolution is apparently broken down when it is claimed that the environmentally influenced ontogenetic responses of scientists (e.g. their beliefs) are adaptive, beneficial and directed toward the solution of the problem they face. In this respect Hull is fully committed to Campbell's model.

First of all, in line with Campbell, Hull denies that intellectual variation is foresighted. Scientists have no "knowledge" that their hypotheses are going to be selected, This answer is sufficient to show that the sense in which variation is guided in conceptual evolution has nothing to do with Lamarckism if by this we mean the claim that scientists "know" (as Lamarck's organisms did) what they need in order to solve their environmental problems. In this sense the Lamarckian claim implies the presence of foresight, which is not involved in scientific evolution.

Furthermore, Hull claims that the notions of tentativeness and unjustifiedness to which, respectively, Popper and Campbell refer, do not merely refer to absence of clairvoyance or foresight, but that they rather refer to the fact that in conceptual evolution there is a link between occurrence of intellectual variation and adaptiveness. The existence of this link is accountable for purely in terms of preliminary selection (e.g. vicarious), and chance. For this reason, Hull (1988a p. 456) points out that the appearance of guided intellectual variation depends on the fact that variants have to pre-fit an already existent, hopefully largely coherent, system of beliefs.

Critics of EET generally argue, from a purely a priori standpoint, that "it cannot be true" that the psychology of discovery can be explained in selective terms. One such dismissive argument is used by Ruse to argue against Hull's EET. Ruse (1989 pp. 208-211) claims that it cannot be sensibly sustained that Gould and Eldredge came up with the hypothesis of punctuated equilibria simply by "stabbing in the dark". As I tried to show in section 3.4, this is an utterly complete misinterpretation of Campbell's model, which does not say that the minds of Gould and Eldredge are chance-like, and that the formulation of their theory was just

[117] D. Hull (1988b) p. 145

due to a lucky accident. Ruse does not take into account conscious and unconscious processes of vicarious selection. He does not take into account that in the time between hypothesis elaboration and formulation, a vast array of cognitive and non-cognitive processes can in principle account for the generative process. Gould and Eldredge knew they had to propose a "scientific" theory, hence some structural requirements concerning what can be acceptably seen by the community as a scientific theory had to be respected. These requirements act as selective constraints in Campbell's model, analogously to biological constraints limiting the optimising power of natural selection. For instance, Ruse does not take into account that the gradualist thesis had been subjected to criticism since Darwin's time, and that knowledge of such criticism inevitably played a selective role, leading to the rejection of those untenable hypotheses that were deemed irrelevant to solve their problem, thus reducing design exploration. Ruse does not take into account that it might be possible that Gould and Eldredge came to formulate this hypothesis after much conscious and unconscious memetic recombination, wild speculation, guessing, etc. He also does not take into account that error and luck might have had their part too, as error in replication has a fundamental part in genetic mutation. Campbell's model potentially explains all these features of the emergence of the theory of punctuated equilibria, despite Ruse's unarticulated objection.

One important epistemological implication in endorsing Campbell's model is that EET narrows the gap between considerations pertaining to the context of discovery and those pertaining to the context of justification. This is because an account of the process of scientific discovery already involves issues that pertain to scientific selection. This means that EET's normative endeavour includes issues traditionally seen as pertaining to the context of discovery (cf. section 5.2.3).

It must be noted that the emphasis on the variational aspect of the scientific process does not solely pertain to EET. Feyerabend's recommendations to think counterinductively and to hypothesise profusely are meant as norms of science. That is, as recipes to increase the rate of change (and, a fortiori, progress). In the same sense, Campbell argues that scientists "should" engage on wild speculation (e.g. in the form of generation of unlikely or even impossible views, or in the form of trying out all the conceivable combination of ideas), that they "should" force themselves to engage in "unbiased" generation of ideas (where the bias is given by the amount of knowledge already accepted), and that they "should" try to articulate the implications of what is already known (i.e. by deductive means, which is clearly a selective process). For Hull, retention is not seen as a bias that "should" be eliminated or whose influence "should" be reduced at all costs. In contrast, Feyerabend argued that if epistemic agents remove all biases affecting the production of intellectual variation that are imposed on hypothesis-generation by already accepted knowledge, then progress is more likely to ensue. On the contrary, Hull and Campbell stress that the success and cumulativity of science are based on the possibility of knowledge retention. As usual with adaptive systems, the problem is about fitness maximisation, which requires, depending on the conditions of the system, different kinds of equilibrium between variation and retention.

4.6 Science as a social and intentional selective process

Critics of EET claim that the selection of hypotheses in science requires intentionality, purpose and goal-directed behaviour on the part of scientists. For this reason they say that selective explanations do not satisfy the third condition for explanatoriness presented (cf. start of section 4.5). This means that the explanatory force of EET's selective explanations is reduced, because the success of science would not depend on scientific selection *qua* selection, but on the mode in which scientific selection operates in science, namely, its intentional nature. Natural and scientific selection are therefore fundamentally different processes. As a consequence, EET is slave to a theory of intentional behaviour (Amundson 1989). In this section I shall assess Hull's arguments to debunk the significance of this argument. In order to show this Hull proposes to see the social adaptive system of science as an invisible hand process like natural selection.

Assessing Hull's arguments against intentional explanations is difficult because he seems to make contradictory claims. On one side, he does not deny that scientific selection is an intentional process, and he agrees that scientists solve problems because they aim to. On the other, he claims that "Some meme-based selection is intentional; most is not" (1988b p. 146) and that "Although science is as intentional as an activity can get, the effects of this intentionality are minimal" (1988a p. 474). The latter sentence epitomises the apparent contradictory nature of Hull's explanatory strategy. I have now to interpret these claims.

One of the claims Hull makes is that a particular kind of intentional action is not essential or necessary for science to succeed in achieving its institutional aims (i.e. objectivity and truth). Hull rejects a particular view of science according to which scientists' behaviour can be retrospectively reconstructed as to involve just intentional striving to seek true knowledge. Hull rejects this picture of scientific selection because scientific practice shows that scientists intentionally strive to pursue a multifarious array of personal and selfish motivations. By means of this claim what Hull rejects is, in other words, not the thesis that scientific selection is an intentional process, but that a certain aim-directed intentional action is essential. On this first claim depends an alternative explanation about the way in which science's success ensues: individual scientists and other scientific interactors (e.g. research groups) intentionally try to achieve their selfish aims; intentional selfish behaviour on their part causes the achievement of the institutional goals of science (i.e. consensus formation and, perhaps, objectivity). The important feature of this explanation is that it is not assumed that the institutional goals of science and the goals of scientists perfectly coincide. Even though we have been brainwashed to think of the creature scientist as first of all a seeker of truth, Hull contends that such an hypothesis does not stand empirical scrutiny, that it is descriptively inaccurate. Scientists are not superior creatures sacrificing individual benefit and the fulfilment of selfish goals in order to achieve the greater good (truth). Rather, in science often individual interest coincides with duty. Science is a tribal and self-policing system in which virtue (i.e. institutional goal of science) and individual benefit (i.e. enhancement of conceptual inclusive fitness and all connected benefits - rise in status, authority etc.) go hand in hand. The reward system in science is organised in such a way that the more self-serving motivations tend to have the same effect as more altruistic motivations. As a consequence, success is an outcome of the cooperative and competitive nature of the organisational structure of science. Note that this claim is compatible with the view that scientists intentionally select theories because they want to achieve a particular aim.

However, the important thesis Hull presents is much stronger, and involves also the rejection of the relevance of the selfish intentional action considered above. Sometimes Hull argues that selfish intentionality is not necessary for science to succeed:

"It is not absolutely necessary because sometimes scientists have made what turn out to be great advances quite accidentally. Chance certainly favors a prepared mind, but a scientific advance is no less of an advance because the problem which a scientist happens to solve was not the one he or she had intended to solve."[118]

This claim means that an action planned to achieve a particular aim can have a different effect, and that for this reason accidentality plays a major role in the explanation of science's success. In this respect, what Hull shows is not that intentionality is not involved, but, rather, that intentionally aiming at y and achieving x is not considered intentional activity. This seems to me a confusing claim because misdirected behaviour remains intentional.

At other times Hull argues that the effects of intentional behaviour on the possibility of science to succeed are minimal, since intentionality at best influences the tempo and efficiency of directional conceptual change, but it cannot be responsible for science being (locally or globally) progressive. Hull argues that science looks directional and progressive because of our retrospective biases. If we edit the history of science enough, then the influence of intentionality is great. Such editing makes science appear effective in achieving its aims, but a careful look at practice instead shows that failure to solve problems is endemic, that success is rare, and that chance has a major role to play. In order to substantiate this thesis, Hull

[118] ibid. p. 147

proposes an argument from analogy: as in ethics intentional action is crucial, so it is supposed to be in science; however in moral matters we have not witnessed any kind of progress, unlike in science; therefore, the causes of the different kind of directionality must lie somewhere else (1988a p. 474). Ambiguously, Hull states that:

"If everything about the natural world were in a state of haphazard flux, scientific theories would also continue to change indefinitely, not just because scientists continue to change their minds about nature but because nature itself is changing. Goal-directed behaviour can have a direction in a global sense only when the goal stays put......Conceptual evolution, especially in science, is both locally and globally progressive, not simply because scientists are conscious agents, not simply because they are striving to reach both local and global goals, but because these goals exist. If scientists did not strive to formulate laws of nature, they would discover them only by happy accident, but if these eternal, immutable regularities did not exist, any belief a scientist might have that he or she had discovered one would be illusory."[119]

I think this argument is self-defeating. Hull seems to deduce from the trivial fact that intentionality is insufficient to explain progress, that it is unnecessary most of the times. What Hull only shows is that intentionality is not a sufficient explanation of science's success. But this does not mean that it is not a causal factor affecting such success. Thus, the above arguments proposed by Hull do not show that scientific selection is not intentional.

The best argument proposed by Hull to show that intentionality is irrelevant in science concerns the explanatory strategy he endorses. Hull argues that scientists are most of the time not conscious that they strive for a particular selfish motivation. Even though awareness is not essential for intentional action (e.g. organisms that show intentional behaviour might be "forced" to act in apparently goal-directed ways by genes or other adaptive constructs), lack of awareness is detrimental to the undertaking of finding out the real motivations of scientists. And this will affect the possibility of formulating generalisations about the influence of "kinds" of personal motivations. As a consequence, reference to intentional behaviour, given the multifariousness and idiosyncratic nature of scientists' personal motivations, is

---

[119] ibid. p. 147

unexplanatory. Hull's EET aims at capturing generalisations about the scientific selection process. In order to do so, the relevant causal factors must be identified. EET's generalisations, so Hull argues, cannot concern idiosyncratic intentional behaviour, but invariant features of the scientific process. Such features pertain to the social organisation of science. For this reason, the explanation of the success of science does not depend on intentional action, but on the mechanism of selection Hull posits.

Hull's alternative explanation of success is couched in selective terms. He tries to reduce the role of intentionality in science by making scientific selection more like natural selection. As in natural selection adaptation emerges without any conscious design on the part of organisms, so in science knowledge growth emerges without any conscious intentional action on the part of scientists. Scientific selection becomes, like natural selection, an invisible-hand process. The social organisation of science, the mechanism of selection based on credit, curiosity and check that Hull posits, is like natural selection:

"As science is now practiced, it is a combination of planned and unplanned, conscious and unconscious decisions, *but it is the invisible hand that tends to keep scientists on the straight and narrow.*"[120] (my italics)

The notions of invisible hand and "unintended" consequence are central in Hull's EET. Hull proposes an invisible hand explanation of the success of science in which some kind of order emerges in the whole from the dispersed interaction of the parts. The first feature of invisible-hand explanations is that intentional design is not involved (Cosmides & Tooby 1994). The second feature is that an explanation of some emergent phenomena is provided in terms of the interaction of lower-level entities which produces a higher-level outcome.[121] These explanations require the division of cognitive labour between lower level entities. The macro-level outcome is the result of lower level entities' actions that are not directed towards the realisation of the higher-level outcome. No intentionality is (necessarily) involved on the part

[120] D. Hull (1997) p. 125
[121] The higher-level outcome is not necessarily beneficial. In Hull's case it is by assumption.

of lower level entities to come up with the emergent pattern. By using invisible-hand explanations, Hull is literally proposing a theory of emergence: the phenomenon of the growth of knowledge emerges because of the organisational structure of science as it is practised. Somehow, the institutional goals of science "emerge" because of the complex interactions between the entities constituting the social system of science; entities at a lower level (e.g. individual scientists) act in such disparate and complex ways (only accountable in populational terms) as to bring about a fit between our knowledge and nature. Selective explanations of this kind do not require postulations concerning properties of individuals (e.g. intentionality), but they postulate the causal efficacy of a selection process.

The question now becomes whether there are any reasons to prefer a selective to an intentional explanation of the success of science.[122] In my opinion such a question depends on what we want to explain. In order to show this, let us consider Sober's (1994c) example from the economic theory of the firm. To explain why businesses are profit-maximisers, we can pursue two investigative strategies. First, we can say that successful entrepreneurs are intentional agents who adjust their behaviour to the market conditions in order to reach their aims. Second, we can point out to the selective mechanism that weeds out inefficient firms by leaving only profit-maximiser ones. Hull's point is that an analysis of the first kind is not as good as an analysis of the second kind as far as the success of science is concerned. To see why, consider that in the first case we explain the success of firms A, B and C (etc.) by referring to the individual properties of their owners. However, the individual actions that led to their survival might not provide a basis for generalisations concerning the optimal actions that entrepreneurs should perform to achieve profit-making. This is because these token-level individual actions might not be subsumable to a particular kind of type-level behaviour. For instance, the owner of A can achieve profit by being a mafia-like entrepreneur (e.g. threatening both workers and rivals). Owner of B can instead be far-sighted, investing large amounts of money or technological and market-product research. Owner of C can be a good team builder,

[122] Invisible hand explanations belong to the more general class of selective explanations.

motivating his staff by letting them participate in the sharing of profit. In all these cases it seems hard to find anything in common at the level of individual behavioural analysis between the processes that led A, B and C to survival and economic success. If the only way to generalise about all these kinds of token-level behaviours is by describing such behaviour as "intentional", then it could be contended that this generalisation is not useful. This is the conclusion Hull reaches as far as the role individual behaviour in science has for the explanation of the success of science. Scientists might intentionally aim at some goal, but no useful generalisations can be found concerning the multifarious, idiosyncratic and selfish behaviour of individual scientists.

To illustrate Hull's argument, let us consider yet again his main test case. Hull explains the success of cladistics over rival systematic philosophies by showing that phenetics was weeded out. The emergence of community consensus as far as the success of cladistics is concerned can be only properly explained, Hull argues, by pointing to the effects of the demic structure of science. Cladists were a more cohesive, tight, and orthodox group than pheneticists. The tight professional alliances they formed causally influenced the dissemination of their memes. Cladists were, in brief, much more efficient in pursuing their group aims. As a consequence, pheneticists were selected against. The ultimate effect is that pheneticists' memes and hypotheses did not spread as much as those of cladists. Hull contends that this pattern of evolution (i.e. consensus formation) is instantiated many times in the history of science. Therefore, the social structure of science (in this case particularly its demic structure) is highly relevant in explaining events of this kind. The selective explanation for the success of cladistics Hull presents does not refer to cladists' intentional action directed towards the achievement of their personal and selfish aims. For Hull what is crucial is that success is an unintended effect of individual scientists' selfish intentional behaviour. What Hull presents is an hypothesis about how cultural group properties were selected, and how this selection led to the success of cladistics. The mechanism of selection is at the group level. The important aspect of Hull's explanation is that group properties are causally more important than individual properties (i.e. intentionality), because it was the tightness of the cladists as a group that led to their success. Hull's claim is therefore about causal primacy: the mechanism he posits is the most important cause of science's success. In this respect I do not agree with Hull (see below). But before turning to criticise the causal-primacy claim, let me stress two important points.

In my opinion, there do not exist general or a priori reasons to prefer invisible hand over intentional explanations, because both explain a particular aspect of the issue. If we are interested in the ways in which individual scientists should act so as to maximise the achievement of their cognitive aims, then an intentional explanation might be preferred. If we are interested in the mechanisms that govern the social aspect of knowledge acquisition (its populational aspect), then selective explanations seem to be more relevant, since only these point to social features of the scientific process that have explanatory relevance. Hull should not conclude that only the second explanatory approach is valuable.[123]

Hull's argument against intentionality is better seen as an argument against individualistic analyses of the success of science. In this sense, Hull extrapolates a fallacious general thesis (i.e. irrelevance of individualistic epistemology) from a valuable point concerning explanatory strategies (i.e. rejection of methodological individualism). One of Hull's aims is to provide an alternative to methodological individualism in epistemology, which cannot account for the causal relevance of group and social properties, since they are considered to be mere expository conveniences reducible to individualistic terms. For Hull, such properties are instead causally efficacious, and help in explaining, for instance, cladism's success. Hull's EET is a populational theory of science which eschews reference to properties of individual

[123] One of the reasons why Hull reaches this conclusion depends on his reluctance to view biological and social systems as "intrinsically" different. Some social theorists contend (e.g. J. Elster - "Ulysses and the Sirens") that while biological systems are to be explained in functional terms, social systems cannot be. As a result, the theorists continue, the explanation of the workings of social systems must be couched in intentional terms. Hull rejects this negative thesis and tries to apply a functionalist analysis to both kinds of systems. Even though I am sympathetic to Hull's endorsement of functionalism as a general explanatory framework, I do not believe that intentionality should be discarded tout court. Rather, a naturalistic interpretation of intentional behaviour compatible with functionalism should be sought.

scientists, but that tries to capture invariances at the level of populations. I believe that this way of looking at science is fruitful because it discloses neglected aspects of the scientific selective process. What I dispute is that, even though populationism provides a new way to understand historical phenomena, it does not imply that intentional behaviour is irrelevant, nor does it imply that all individual level behaviour is explanatorily irrelevant just because we can tell part of the causal story from the populational perspective.

In Hull's case, populationism seems to be not purely a methodological perspective (i.e. the antithesis of methodological individualism), but an ontological one, namely, that individual level properties can be reduced to populational ones (i.e. radical holism). However, a better alternative to methodological individualism than radical holism is anti-reductionism, that is, the thesis according to which the success of science must be couched both in populational and individualistic terms, all of them in principle reducible to selective ones, and therefore compatible with EET as a general theory of selection processes. Hull seems to re-propose Putnam's (1975) peg-hole argument by arguing that intentional explanations give irrelevant information in order to explain the success of cladism, which can be more accurately explained in macro-level terms by referring to group properties. This argument is fallacious like Putnam's peg-hole argument is, and like any a priori antireductionist argument put forward to defend the explanatory irrelevance of lower level properties is. Radical holism not only rejects the reducibility of higher-level processes to lower-level ones, but it also rejects the causal efficacy of lower-level mechanisms, which are explained away by pointing out that they are mere consequences, effects, artefacts of the causal action of mechanisms of a higher-level order. I believe that radical holism is not an option open to EET. The open alternative is anti-reductionism. As some philosophers argue (Sober et al. 1992 p. 115), "....the issue is not whether the individual level analysis can be eliminated, but how it should be linked to macro-level social analysis." Endorsing anti-reductionism means that both intentional and invisible-hand explanations are necessary to explain the success of science. In Hull's test case, an antireductionist outlook would refer at least to two causal factors. First, intentional explanations refer to the individual behaviour of cladists (e.g. their aggressiveness and singleminded conviction in defending the theoretical tenets of their view) that led to the spreading of the memes of cladistics. Second, Hull's invisible hand explanation focuses on the group properties that led to cladism's success (e.g. its cohesiveness). In this particular case, the two kinds of explanation mutually provide a fuller account of cladism's success. They are compatible in this case, even though I do not claim that invisible hand and intentional explanations must always be compatible.

I argue that we can give explanations of the success of science from at least two levels of analysis. The first level concerns individual scientists and intentional action, the second concerns the population level and the social structure of science. Hull's epistemological analysis regards merely the population level, the only level, he claims, at which fruitful generalisations concerning science as a process can be found. I contend that Hull's populational explanations of the success of science only identify part of the causal factors influencing science's success.

As a result, Hull's thesis that his invisible hand has causal primacy in explanations of epistemic fit seems to me unsustainable. Issues about causal primacy are empirical. Even a brief look at the history of science shows us that different episodes in science are causally affected by different causal factors. Eliminating intentional behaviour completely cannot provide a sound explanatory strategy. As a matter of fact, Hull does not show that the influence of his mechanism is even primary in his test case (cf. section 4.8). Therefore, generalising its influence to the whole of science seems to me even more dubious. For this reason, explanatory pluralism is a better investigative strategy than invisible-hand monism:

"In the case of a theoretical account of science, on the other hand, the intentional capacities of scientists suggest that a variety of different processes are plausible. All or most scientists may, indeed, seek to maximize the credit they receive for their contributions. But they may surely have many other goals besides, some of which will have rather different effects on the progress of science. Though scientific intentionality is no objection to the possible importance of Hull's mechanism, it is an objection to its uniqueness."[124]

In my opinion, the invisible hand is just part of a whole series of selective processes operating in science at many levels involving individual scientists, research groups and even the scientific community as a whole. By claiming that his invisible hand is the "unique" analogue to natural selection, Hull constructs his EET as based on a strong interpretation of the analogy between biology and science (cf. section 4.7). But this is a simplification that comes at a cost, namely that Hull's EET completely rejects the relevance of the selective processes operating at the individual level.

Ultimately, I think that Hull's argument against intentionality is misdirected. The reason is that he mistakenly thinks that only if scientific selection is rendered, like natural selection, an invisible hand process, then EET can be saved. Hull takes too seriously the criticism that intentionality is a fundamental difference between scientific and natural selection. As a matter of fact, I believe that scientific selection is not entirely analogous to natural selection, and also that this fact is irrelevant as far as EET's epistemic value is concerned, especially given that selection theory might provide a satisfactory theory of intentional behaviour. In this sense, I cannot see why invisible hand explanations cannot be coupled with intentional ones in order to achieve more complete explanations of particular episodes in science. Such coupling would, in my opinion, provide a fuller explanation of the success of cladism (cf. section 4.8).

The analogical agenda plays too big a role in Hull's EET (cf. section 4.7). Hull seems to argue on the following lines: as natural selection only explains how a population evolves regarding the frequency of a certain trait, so scientific selection only explains how a population of scientists evolves regarding the frequency of a certain meme or hypothesis. The important point is that in both cases selection does not explain why an organism has a certain phenotype, that is, in the scientific case selection does not explain how a certain scientist

[124] J. Dupre' (1990) p.75

holds a certain belief, meme, hypothesis (e.g. because she makes a certain selective choice). This view is criticisable for reasons similar to those that have been used in the philosophy of biology (Neander 1988). And for the additional reason that a populational history of science would have many deficiencies. Often what is important to ascertain is how a certain scientist came to hold a belief, what was the cognitive process that led her to figure out that a certain meme had to be endorsed. Such a cognitive process can be intentional or unintentional, but Hull's EET is silent about its nature.[125]

Scientific evolution is partly the outcome of an invisible hand process in which misdirected (but still intentional) actions of individuals cause the system to realise its goals. The issue is not, therefore, that scientists' behaviour is not goal-directed, or that scientists often do not intentionally choose the hypotheses they accept. The issue is rather that science can be seen from a different perspective. From this populational perpective the explanation of the growth of knowledge is due to the social features of the structure of science (e.g. the demic structure of science). Hull's populational explanations do not refer to the vagaries of intentional behaviour apart from understanding such behaviour as directed towards the maximisation of conceptual fitness. In this sense, Hull tries to show how the good of individual scientists coincides most of the time with the good of science, that is, how the motivational goals of scientists, when properly constrained by the workings of the credit-checking system, by mechanisms of reward and punishment that in science act as to curb individual and group cognitive biases, lead to the achievement of the institutional goals of science (e.g. consensus). Science can certainly be seen as an adaptive and largely adapted self-policing system. I agree with this completely. But this picture of the selection process must not be coupled with the minimisation of the role of intentional behaviour in science, by contending that such an intentional (and, I believe, selective) process is just reducible and accountable for solely in terms of unintended achievement of science's institutional aims. That science has a particular

[125] I believe that Hull's populational approach can in principle be accompanied by an ontogenetic selectionist account of the intentional behaviour of individual scientists that would obviate such a limit.

social organisation whose mechanisms partly determine science's success is not a substitute for this kind of intentional striving.

A further limit of Hull's EET is that he offers only scattered comments concerning the nature and causal role of the institutional norms or ideological commitments that regulate and direct scientists' behaviour. This is rather paradoxical for an EET, as Hull's is, that is solely a social epistemology. The teleological process of science is governed by specific norms and by specific goals that exist independently of the multifarious goals of scientists. In this respect, Hull's views would be directly falsified if it were demonstrated that scientists behaviour is basically motivated by one ultimate aim (e.g. truth), despite the variety of the proximate and selfish aims they tend to achieve. That this is not the case, Hull claims, is shown by the history of science. Even if we grant this descriptive thesis, the fact remains that Hull's EET does not tell us enough about the norms governing scientific behaviour. EET should not be confined to account for the teleological nature of the scientific process by relying on Hull's concept of "unintended consequence", since the social norms and the self-policing mechanisms of science have a causal efficacy in directing the selfish intentional action of individual scientists. For instance, Hull argues that if there were no such norms at the heart of the social practice of science, science would not progress. The institutional norms of science provide at least a motivational foundation for science:

"The motivations of individual scientists are important, especially in histories that follow the careers of individual scientists. They are also important because of the incredibly complex relation between individual motivations and institutional norms. Certainly groups are more than sets of people. They are people organized in particular ways. The issue is how such organization can influence the behavior of the agents so organized. Can norms function as causal agents? Can entities at all levels of organization be understood solely in terms of their constituent elements and the relations between them?"[126]

Hull does not say whether such norms are Popperian regulative ideas, or whether they are more like Campbell's anti-tribal norms of science, or if such norms of science should be rather

[126] D. Hull (1988a) p. 393

seen as Kuhnian ideological commitments. What is clear from the quotation is that such norms might have a causal role. This would provide an additional argument against methodological individualism, since the success of science would depend on factors that are not completely accountable for in terms of scientists' behaviour (cf. section 5.1).

In this section I have shown that Hull's arguments directed at showing that intentionality is irrelevant in science are misguided for many reasons. I now pass to consider a basic defect of Hull's science of science.

4.7 The analogical agenda

Hull argues that if the aim of the science of science is to find regularities about the process of science, these must be found at the level of generality at which they are found in biological evolution. The other, but less general, aim of Hull's EET is to provide a science of science that can function as a normative social epistemology. Thus, he argues, the only regularities that can be found in science concern general features of the social organisation of science. I believe that these two aims are at odds with each other in Hull's EET. More specifically, I believe that Hull fails to achieve the most general aim of singling out interesting nomic generalisations about the scientific process in terms of the replicator/interactor/lineage categorisation of selective processes. In particular, I have not found any link between the sociological generalisations he provides, and the more general laws of selection, apart from purely analogical claims pertaining, for instance, to the similar function that the demic structure plays both in biological and scientific evolution. Hull does not deliver in his major aim to tie generalisations about how science is organised, structured and practised in the West at this time, with more general laws of selection. Ultimately, it seems to me that Hull's categorisation and his notion of conceptual fitness are largely irrelevant as far as his social epistemology is concerned.

One basic reason why Hull fails in this endeavour depends on the scope of his project. Hull wants to achieve nomological aims (finding "laws" of selection processes), unificatory aims (providing a general theory of selection processes that is applicable to any selective process - be it biological, cultural, or scientific in origin - by adopting the same terminological and theoretical framework), as well as developing a sound science of science. This project is majestic. Even though Hull's EET grasps sociological generalisations about the scientific process that can be deemed to be interesting and fruitful for purposes of epistemological analysis, this fact does not imply that such generalisations are instantiated as more general laws of population memetics or as even more general laws of selection.

In the previous section I also pointed out that Hull's endeavour is vitiated by strenuously defending the thesis that natural and scientific selection are not fundamentally different processes. I consider this task as irrelevant, simply because the selective processes operative in science are different from those operative in organic evolution.[127] As a consequence, if scientific selection is intentional, this does not affect the value of the sociological generalisations Hull provides. We must be careful to discriminate between Hull's EET as a theory of science, and the vastly more general goal of providing a very general selection theory. While the first objective is partially fulfilled by Hull, the second is certainly not. In this sense, all the controversy about the possibility of a general theory of selection processes framed in terms of the replicator-interactor terminology seems to me largely irrelevant as far as an assessment of EET is concerned.

The problem is that the analogical agenda infects Hull's epistemological analysis. In this respect, one of the problems of Hull's view is that the concept of conceptual fitness might be vacuous. Hull uses this concept because it provides the only way, in his opinion, to generalise about the idiosyncratic intentional behaviour of individual scientists (thus allowing to eliminate reference to intentionality, saving the strong reading of the analogy between natural

[127] Having said this, EET is still committed to the view that the variation-selection model is a valid model for understanding any evolutionary phenomena. Only in this general sense, I believe, natural and scientific selection are not fundamentally different processes (cf. section 5.2.2).

and scientific selection - cf. section 4.6). Any token-level intentional behaviour is grouped under the type-level category of conceptual fitness enhancement. Scientist A behaves so as to achieve aim A. Scientist B so as to achieve aim B, and so on. What all these aim-conducive behaviours have in common, Hull claims, can be captured by the notion of conceptual fitness, because scientists are assumed to act as to maximise this measure of fitness. In this sense, Hull contends, the concept is analogous to the concept of biological fitness. As biological fitness is multiply realised by indefinitely many physical properties, so conceptual fitness is multiply realised in indefinitely many cognitive ways.[128] In biology fitness is supervenient on the physical properties of the unit of interaction and on the environmental properties the interactor inhabits. The situation seems entirely parallel in science. Fitness is a property of individuals and communities of scientists. In all these cases fitness is supervenient on cognitive (physical) properties of scientists (e.g. their beliefs) and environmental properties (the properties of the scientific community the scientist inhabits - e.g. various kinds of metaphysical, methodological and theoretical hypotheses). In both cases fitness is a placeholder for its supervenient properties. The biological notion is useful because it allows biologists, among other things, to single out the common features of patterns of events, abstracting from the idiosyncratic physical characterisation of the fitness property in different organisms and species. Can the same be done for epistemology? I doubt it.

One reason is that what is mainly relevant epistemologically is to pinpoint the physical and environmental features of the selective process in detail, rather than providing generalisations concerning the relative conceptual fitness of scientists and research groups. To claim that fitter scientists and fitter groups causally affect the spreading and dissemination of ideas is a platitude that does not add any content to Hull's sociological generalisations. For instance, to claim that the success of cladism as an interactor caused the differential perpetuation of the memes (e.g. experimental techniques, methods and hypotheses) of cladistics because cladists

[128] The first problem of this idea is that even intentionality can be multiply realised, so that it would serve the same function as conceptual fitness. Furthermore, Hull assumes that intentionality is a property of individuals, instead of stretching the notion as to include, for example, social groups.

as a group were tighter, more orthodox etc. is one thing. To add that cladists were, as a group, conceptually fitter than rival groups is just redundant information.

Hull claims that the source of epistemic fit is accountable by referring to the interactions of scientists among themselves and with the world, and that such interactive procedure is a selective process, and that, consistently with EET's perspective, selection is the only source of fit and consensus formation. More specifically, Hull argues that social mechanisms of selection are mainly responsible for the explanation of the success of cladism, but such a social epistemological explanation is distinct from the evolutionary analogue couched in terms of the notion of conceptual fitness. In brief, I believe that Hull's sociological and evolutionary generalisations are two different things not sufficiently linked in his EET. While the first are explanatory (because they refer to the social mechanism causally affecting the scientific selection process), the second provides mere analogies and platitudes about the nature of the selective process of science. This is because what epistemologists and historians of science want to know are the specific causes of consensus formation, which cannot be put all together under the summary notion of conceptual fitness.

Among the causes of consensus formation a variety of factors might be included, ranging from those identified by internalists and externalists, to those concerning the social structure of science. What I contend is that by finding analogous versions of concepts that have found fruitful application in biology, Hull pays the price of being unable to account for the different nature of the causes that determine the success of scientists and, more importantly, the success of the ideas and hypotheses they propose. In particular, Hull's evolutionary generalisations in terms of conceptual fitness seem to me not only compatible with the genuine and innovative sociological ones he proposes, but also, paradoxically, with the intentional, internalist and externalist explanations of the success of science he wants to replace.

For instance, Hull's evolutionary explanation of the success of cladism framed in terms of the conceptual fitness concept is compatible both with the internalist explanation that stresses the empirical adequacy and theoretical superiority of cladism's memes, and with Hull's sociological explanation. Internalists could argue that the success of cladism can be explained by pointing to the vices of phenetics. For instance, pheneticism was based on the highly criticisable methodological requirement according to which the important distinction between homologies and homoplasies is of no interest to taxonomy, since all that counts are judgements of overall similarity. In turn, the notion of overall similarity has been strongly criticised because it does not seem to refer to an objective feature of reality, especially given that pheneticists disagree on the method of characterising what overall similarity is. Pheneticists also endorsed a very narrow empiricist taxonomical methodology dictating to look only for "pure" observations (i.e. uninfluenced by what they thought was always unreliable evolutionary information). This notion of observational purity is criticisable for reasons philosophers of science know too well. Hull's explanation points to a different series of causal factors, ranging from the tighter professional alliances between self-professed cladists, to their group efforts in establishing editorial control over academic journals and institutions. These two kinds of explanations are both compatible with those merely analogical and vacuously evolutionary explanations in terms of conceptual fitness, since in both cases it is clear that cladists were conceptually fitter than pheneticists.

My point can be further illustrated by considering Hull's explanation of why stealing is less frequent than fabrication in science (cf. section 4.3). The sociological explanation Hull offers refers to the fact that, since fabrication, unlike stealing, hurts potentially the whole scientific community rather than just one or few interactors, there exist social mechanisms that act as to minimise fabrication by punishing scientists who completely fabricate data. Other times, Hull adds that, seen from the conceptual fitness perspective, lying is punished more severely than stealing because it will induce a decrease of absolute average conceptual fitness in the community, rather than simply affecting the conceptual fitness of one or few scientists. I contend that the sociological explanation is interesting, while the second might turn out to be purely analogical and unexplanatory. In brief, nothing is added by reference to the putative notion of absolute conceptual fitness.

In a nutshell, these illustrations might prove that the notion of conceptual fitness is epistemologically irrelevant, since to say that fitter scientists are more successful (and that success has a causal role in explaining the dissemination of their ideas and hypotheses) does not explain what makes scientists select memes for, and does not differentiate between the series of causes of scientists' success.

Hull's analogical agenda sits uncomfortably with his social epistemology. What the former side of Hull's EET offers is an analysis of science that, instead of looking at the similarities (many) between the process of biological and scientific evolution, should also focus on the dissimilarities. One such dissimilarity consists in the different explanatory role the concept of fitness has in biology and science. Hull's EET is a useful framework to understand the causal role the social mechanism he posits has on scientific selection. On the contrary, it provides only trivial generalisations when evolutionary analogies are pushed too far. I shall now move to assess the genuine and original aspect of Hull's EET.

4.8 EET as a social epistemology

In this section I shall focus on the social character of Hull's EET, instead of its analogical one. I will try to answer some questions about the significance of Hull's EET, particularly concerning the value of his sociological generalisations and explanations. The first question regards the completeness of Hull's EET, and its role as alternative to other epistemological approaches. The second pertains to Hull's contention that we can extrapolate generalisations concerning the influence of some social features of the scientific process.

I generally agree with Hull that reference to neglected social factors about the organisational structure of science is epistemologically relevant, but I disagree with the contention that his mechanism of selection is the causally primary factor affecting the success of science. This is because I doubt that EET is either a complete epistemology, or "the best" epistemological approach. To illustrate my thesis (cf. also sections 5.2.3 and 5.3), we need to understand properly what EET offers to epistemological analysis. Let us start from the epistemological analysis of singular events, of particular episodes in the history of science. Does EET provide a historiographic model? It does, but only, in my opinion, in the sense that EET provides an additional tool to extract information about the social aspect of the process of knowledge growth. This interpretative tool does not generally yield "better" information about the event to be understood. It is, in this sense, merely an alternative tool to those offered by internalist and externalist epistemologies. Issues about causal primacy are highly contextual, depending on the nature of the event that needs explaining. To ascertain which causal factors are operative in a specific episode is eminently an empirical question. Different episodes can be affected by many kinds of factors whose causal role is highly variable, and whose relative importance is context-dependent. At least this is how I interpret the evidence coming from science studies, and from my limited knowledge of the history of science. The basic point is that it is difficult to generalise a priori, assuming that certain causal factors must have a preponderant role in every case. Furthermore, it is generally a multiplicity of causal factors that influences the particular event to be explained. This seems to be the case in Hull's chosen main test case. Hull's explanation of the success of cladism is only part of a complete explanation, additional and compatible with internalist explanations (cf. section 4.7).

It is for these reasons that Hull's arguments put forward in order to debunk the epistemological significance of alternative approaches (cf. 4.2) are unsatisfactory. Hull's contention that his EET provides the best explanations of the success of science cannot be sustained because I doubt that any kind of explanation (couched in intentional, rational, internalist, externalist, populational, social, individualistic, invisible hand terms etc.) is a priori preferable to any other. General arguments of this kind, and general claims about causal primacy, are hard to confirm. And Hull is certainly unsuccessful in this respect.

The moral of my argument is that no epistemological approach is a priori in a better position to explain science's success, given that different historiographic models generally emphasise the influence of different causal factors whose causal influence is case-dependent. This view is in line with the epistemic pluralism I endorse.

For parallel reasons, I believe that Hull cannot sustain his thesis that individualistic epistemologies are not relevant to explain science's success (cf. section 4.6). This is because social and individualistic epistemologies have different programmatic aims. The latter aim at finding descriptively accurate and significant explanations of individual choices and behaviours, with the ultimate task of coming up with normative recommendations for individual scientists concerning how to best behave in scientific contexts under certain conditions and given a certain epistemic aim. Hull's EET, as a social epistemology, has a different agenda. On the descriptive issue, it must show how the social (and hopefully adaptive and adapted) features of the process of scientific selection might influence the growth of knowledge. On the explanatory issue, it must identify the social selective factors that caused consensus emergence. On the normative issue it must mainly give recommendations on how to best organise the practice of science in order to fulfil the aims that the community of scientists deems to be topical. Given such axiological difference, I believe that Hull's invisible hand explanations will not generally substitute, but only complement individualistic explanations (which are not assumed, for similar reasons, to be primary). I do not doubt that such explanations might be extended to many scientific episodes, as I do not doubt that the invisible hand explanation of the success of cladism Hull provides is novel and explanatory. I only doubt its completeness.

Sometimes Hull claims that the social mechanism of selection he describes causes the fit between our knowledge and reality. At other times, he claims that:

"At the very least, in the absence of the mechanism which I have sketched, science could be likely to proceed at a very leisurely pace."[129]

I prefer to interpret Hull's EET according to the second, less smug, claim. EET is not an alternative, but merely an important part of a more complete epistemology of science.

It might be argued that the real strength of EET lies not in offering novel explanations about particular historical episodes, but rather in explaining patterns of events. In this sense, EET as a social epistemology provides a novel analysis (i.e. populational, cf. section 5.1) regarding the characteristics that many historical episodes have in common. In this sense, EET's main aim should be that of providing generalisations couched in adaptive terms about the scientific process as a whole. In my opinion, Hull's EET fares better in this respect than as an historiographic model. The most original insights Hull offers do not concern his test case (a single historical episode), but the role of the demic structure of science on the scientific process, the influence of the credit system, and the role of the self-policing processes affecting the scientific process as a whole. In this category I add Hull's hypotheses about scientific group selection (treated in section 4.6), and his equilibrium hypotheses concerning the stable and most functional relationship between cooperative and competitive behaviour in science (cf. section 4.3). In this sense, Hull presents an innovative social theory of science, which, however, is still, as a science, in its infancy. Thus, Hull might be excused for delivering less than he promises. I have already pointed out two reasons for why Hull's EET does not deliver. The first is his reliance on the analogical agenda (cf. section 4.7). The second depends on the general complexity of the undertaking, due to the fact that it is difficult to individuate those social factors that have a positive, long-term and constant effect on the scientific process, while at the same time discriminating between those that have, in different circumstances, a positive and a negative effect. As Donoghue puts it:

[129] D. Hull (1988b) p. 254.

"Unpacking a complex functionally organized system [like science] is difficult because cause and effect inter-penetrate so extensively. Many elements in the system seem to be both a cause and an effect, and they function in both capacities somewhere in the analysis."[130]

In my opinion, the fundamental reason why Hull's EET is unsatisfactory pertains to his inability to distinguish between various categories of social factors affecting the process of knowledge acquisition. Some might have causal relevance highly idiosyncratically, affecting consensus formation only in a few actual historical cases. Others might influence the selective process a sufficient number of times, allowing space to find out proper generalisations. There are then those factors that "should" have causal relevance, in the sense that they are demonstrably cognitive conducive. In my opinion, EET should offer an epistemological perspective from which the social structure of science can be seen as providing both assets and liabilities for knowledge growth (cf. section 5.2.2). Assets of the social structure of science are those features of the social process that reduce individual and group cognitive biases, that promote the communal access and the inter-subjectivity of knowledge, and that more specifically render science a different social system from other practices. In this category we can classify the credit or reward system in science, and more in general the self-policing organisation of science as an epistemic practice. The liabilities of the social structure of science are those properties that function in contrast to the assets, that is, by interfering with the relevant cognitive factors that "should" play a selective role. The discriminative task has both explanatory and normative relevance. In fact, EET's aim of extrapolating recommendations in order to improve the objectivity of our knowledge is based on such discrimination: if a social factor constantly acts as an asset, then the recommendation is to protect and amplify its influence on the scientific process (cf. section 5.2.3)

I am not claiming that Hull is completely unsuccessful in this endeavour. For instance, he does not consider as generally cognitive conducive many of the social factors that are idiosyncratically identified by externalists (cf. section 4.2). In this sense, he points out that scientific communities should try to insulate themselves as much as possible from the

---

130 M. Donoghue (1990) p. 471

influence of external factors if consensus has to be reached quickly. Conversely, Hull considers the mutual monitoring at the basis of the self-policing system of science as a maximising factor. The scientific process is based on the workings of many institutional arrangements that force scientists (e.g. by rewarding them) to behave according to some adaptive social norms. Hull also argues, consistently with evolutionary epistemology's tenets, that such institutional arrangements are not specifically devised, but that they are more accurately features of the unintentionally designed system of mutual exploitation of science. The social system of science is adaptive in this sense.

Despite all these virtues, Hull's analysis is deficient for at least two reasons. First, Hull seems to equate consensus emergence and objectivity. For instance, his explanation of the success of cladism refers to the epistemic role played by a particular social factor, that is, the nature of the professional relationships that determined the groups' exceptional cohesiveness. My point is that while reference to professional relationships might be explanatory in this case, its influence cannot be generalised a priori. Even if it is granted that the peculiar kind of allegiance between cladists might have quickened the process of consensus emergence by ultimately causing the success of cladism, we have no reason to think that in other situations such a social factor will not act as a liability. This would be the case even though we might admit that building strong intra-group alliances will quicken the process of consensus emergence. But to argue that the causal role played by professional relationships will inevitably be positive (i.e. epistemically conducive) equates to arguing that it is assumed that the aim of science is reduced to consensus formation instead of genuine knowledge growth. And such an equation is difficult to justify generally, even if we grant that it is justifiable for the test case Hull describes.

Campbell & Paller (1989) argue in a similar fashion. What Hull provides, they contend, is an EET that is modelled too much on the analogy with biological evolution (cf. section 4.7). But any strong parallelism with biological evolution may or may not help to suggest processes of scientific selection that enhance, optimise and maximise epistemic fit. The effect is that the scientific process becomes a competition between competing "ideologies" that can be applied to religious sects and political factions alike, where "popularity" becomes a criterion of truth. More generally, Campbell and Paller argue, Hull risks being seen as a sociological relativist, according to whom scientists "construct" scientific knowledge as organisms shape their ecological niches. This view, Campbell contends, is partly correct, but also deeply unsatisfactory when "facts" are not considered to have any selective role (cf. section 5.4). Campbell's criticism seems to me too strong, since Hull cannot be seen as a sociological relativist or extreme constructivist, given his arguments against externalism (cf. section 4.2).[131] What remains true is that for Hull consensus formation is seen as the exemplification of success and objectivity, while consensus formation is better seen as a proximate rather than the ultimate aim of science.

The other and related reason why Hull's EET is unsatisfactory is that it does not provide a satisfactory solution to the sociological version of the demarcation problem. Note that Hull believes that his EET describes "the" general mechanism of conceptual change, general in the sense that "all" disciplines that can be termed scientific evolve in the way he describes. He thinks that his mechanism of social selection can be used as a principle of demarcation.[132] Dupre' is correct in pointing out that such an argument is equivalent to an attempt to frame the thesis of the unity of science in sociological terms:

".....scientific unity might be discovered as a feature common to the institutions and social processes of science......if there is some way of setting up the institutions of knowledge production so that beliefs are somehow selected exclusively for truth, empirical adequacy, or some such suitably exalted epistemological virtue, then it might be reasonable to call just those institutions set up in this way the sources of genuinely scientific belief."[133]

[131] It seems certainly true that sometimes scientific disputes are solved because of the selective role played by external factors, or that long standing disputes are solved just because opposing scientists simply die off (Planck's principle – cf. section 5.2.3). Externalist science studies might have an explanatory relevance as far as descriptive and explanatory issues are concerned, and have certainly contributed to revolutionise the ways in which epistemological research is pursued, by helping to dismantle the logical empiricist approach based on the a priori evaluation of meta-scientific claims about science. But two issues are relevant at this point: how typical are the scientific episodes chosen by externalists, and what is the normative relevance of these studies? Hull's answer is that such cases are not typical, and that, as a result, externalist studies have no normative relevance.
[132] D. Hull (1988a) p. 285
[133] J. Dupre' (1993) p. 233-4

Hull's thesis is that, at least a sufficient number of times, bad scientific ideas are weeded out from science by the peculiar social selective process he describes. And because of the competitive and cumulative nature of the social practice of science, in the long run the selective process will lead to the retention of those memes (conceptual systems, hypotheses, theoretical models, concepts, methodological, epistemic standards etc.) that are "most-valuable". In a nutshell, scientific progress of some kind ensues.

Here I am not going to dispute whether it is sensible to find a solution to the demarcation problem in any form, be it sociological or methodological.[134] What I shall dispute is that Hull's solution to the problem is unsatisfactory because such a solution presupposes a previous identification of those social selective processes that increase "competence of reference" (Campbell & Paller p. 243), and that justify a belief in the epistemic superiority of science. To repeat my point, Hull does not properly classify the variety of social factors that he considers to be relevant for scientific progress in terms of their explanatory and normative relevance. In this respect, the influence of the credit system as structured in science seems to me an important and causally efficient feature of science (both from an explanatory and normative point of view), while the effect of the professional relationships is not. This is because they might have a positive effect in some cases (e.g. his test case), but their influence seems to me to be too idiosyncratic and not generalisable (and hence unexplanatory), while it has arguably no general positive effect on consensus formation (thus being normatively irrelevant).

The basic limit of Hull's EET can be also illustrated by considering the causal role of the anti-tribal, antiauthoritarian, anti-traditional and anti-metaphysical norms of science, which is not sufficiently stressed in his theory of science. Hull only hints at the fact that such norms have

[134] What remains to be seen is whether a series of selective processes is peculiar to science in general, or whether specific processes are operative in different scientific disciplines but still provide a sufficient basis to discriminate between sciences and pseudo-sciences. The problem is empirical. Hull's suspicion is that science is "a" process of some kind, instead of a series of processes, contravening the externalist's claim that epistemology cannot aim at useful generalisations (cf. sections 4.2 and 5.4).

a selective role as a motivational foundation of scientists' behaviour (cf. section 4.6), without realising that their major role consists, as Campbell claims (1979 pp. 499-502) in minimising the selective role of irrelevant factors in scientific selection (cf. section 5.2.3). In this respect, I strongly agree with Campbell that the social and anti-tribal norms of science differentiate the practice of science from other tribe-like traditions.

I shall now conclude my assessment of Hull's EET by considering its normative value.

4.9 The normative nature of Hull's EET

So far I have mostly tackled the issue of the descriptive and explanatory value of Hull's EET. Even if we grant the descriptive accuracy and the explanatory fruitfulness of Hull's EET, what normative value has such an account? Hull's invisible-hand explanations are descriptive, not normative (Nozick 1994 p. 314). His EET is a descriptive and explanatory theory aiming at capturing nomic regularities concerning the scientific process. Hull argues that these nomic generalisations have a prescriptive value, since they identify necessary conditions for the growth of knowledge. This sense of prescriptivity consists in the possibility of deriving "norms" for scientist's behaviour, and "norms" regarding the most effective social organisation of science (cf. section 4.1).

I believe that Hull's EET has some implications for social theories of the organisation of science, and that advice can be derived for science policy makers. According to Hull, the best way to justify normative claims about science is by putting them into practice. More specifically, Hull argues that we need to convince scientists that they should explicitly adopt some of these meta-claims, that they should behave in the ways prescribed by empirical theories about science, and then see what the effects are. If the result is that science in the particular area grinds to a halt, then the claim is not confirmed and hopefully refuted. If, on the other hand, the normative advice extrapolated by his science of science is correct, then progress will ensue, and EET's normative endeavour will be achieved and vindicated. By convincing scientists that by acting in particular ways they will reap epistemic benefits, or by organising science in a particular way in order to accelerate the pace of progress, meta-claims are tested. Evidence has therefore a role in testing meta-level claims about science (Hull 1992a).

One problem in deriving norms from nomic generalisations is that Hull does not provide many good generalisations of this kind. This is mainly because his largely analogical analysis of the scientific process prevents him from capturing significant evolutionary generalisations (cf. section 4.7), apart from vague claims to the effect that science must be structured in demes in order to accelerate the process of knowledge growth, or that intellectual variation between and within groups is a preliminary condition for intellectual novelty to emerge. I largely agree with such claims, but Hull seems (and honestly should) promise much more. What needs to be done is to specify the environmental conditions under which the demic structure of science will lead to the maximisation of knowledge growth (cf. my example in section 5.2.3), or to specify the conditions under which a particular equilibrium between intellectual variation and conservativeness is functional in order for knowledge growth to ensue.

A second problem regards the testing of such claims. Even granting that Hull's nomic generalisations are satisfactory, and that we can therefore derive normative advice of some form, empirically testing such advice is a very difficult task. For example, Hull claims that, to test the normative claim that scientists should aim at capturing laws of nature, what we have to do is to convince scientists that they should aim at different goals. If by redirecting scientists' behaviour towards the attainment of other goals the result is decreased growth of knowledge, then the normative advice has to be rejected. Critics doubt that testing such claims is possible or fruitful.

A third problem concerns the general conservatism of Hull's normative endeavour. To illustrate this problem, consider that Hull argues that the contemporary structural organisation of science in Western societies is already functioning decently, that it is already suited to achieve its institutional aims. For example, Hull says that every time the self-interest of scientists is redirected toward the attainment of goals other than the enhancement of their conceptual fitness, certain deleterious consequences concerning the effectiveness of the scientific process can be predicted. If this is so, then the advice for science policy makers is to leave science as it is actually organised, precluding the possibility to ameliorate and improve its epistemic conduciveness. Science policy recommendations would therefore consist merely in protecting the contemporary organisational structure of science. In this sense, critics say that Hull defends the status quo. I partly agree with Hull, but he seems to base this claim on the assumption that science is already an almost entirely adapted system. The implication of such a view is this: what Hull (1988b p. 155) describes equates to "nearly" the best possible organisational structure of science.

I contend that all these are generally symptomatic of a view of science as a functional system that is innovative. The testing problem can be solved by relying on historical evidence and by conducting carefully designed controlled social experiments. Given Hull's example, we would need to specify what sorts of aims scientists should adopt instead of seeking natural laws, and then we would either require a historical comparison of achievements between two different ways of organising the scientific process, or empirical evidence coming from social experiments. Such empirical evidence is certainly difficult to obtain, but it is not impossible to come by. Claims about functional organisation are not more easily tested in biology. This said, it might still be wondered whether progress would accelerate if scientists were only interested in intentionally aiming at achieving the goal of truth, instead of being interested in pursuing purely selfish aims. In this sense, forcing the adherence to the anti-tribal norms of science by preaching them more effectively, or by brainwashing science trainees, might be seen as a recommendation that can have positive effects. Hull strongly disagrees with this view, retorting that to force scientists' individual behaviour to respect such norms would have mostly deleterious effects, and that it is not even necessary because institutional arrangements (i.e. adaptive features of the social system) that direct scientists' behaviour are already operative in science. This points to the third problem, namely Hull's defence of the organisational status quo. It must be pointed out that such a defence is not based on the a priori assumption that any product of selection will necessarily have aim-conducive effects (which would amount to a sociological variant of Panglossian adaptationism), since maladaptation is always a possibility, especially considering that the scientific process is affected, for example, by external factors (e.g. industry-driven research funding) that were largely causally inert in what can be termed the "environment of ancestral adaptation" (i.e. the environment in which science emerged three centuries ago). Science will constantly change its nature, like any other adaptive system. Thus Hull's defence of the status quo is not equivalent to the purely negative thesis that a better and alternative social organisation of science is impossible to find, or that the scientific process cannot be improved upon. Rather, what Hull claims is that the adapted system of science is best suited to fulfil its epistemic aims only if the causal influence of interfering factors is minimised. To draw the last analogy with evolutionary biology, Hull's EET states that science would be completely adapted (it would produce epistemic fit) only if scientific selection as characterised in his model were the only operative force affecting the emergence of adaptation, in parallel to what would happen in biological evolution if only natural selection were operative. This is not the case in the actual world, as natural selection is not the only force operative in biological evolution. Therefore, given the importance that science has as an institution in our time, we should try to understand, first, what are the social forces that affect scientific evolution in particular areas, and then we should try to improve its organisation by means of science policy recommendations that can be derived from the functional and evolutionary analysis EET provides. Having said this, it must be admitted that Hull is certainly an optimist about science, contrary to Feyerabend (1975), who provocatively claimed that society must be defended from the new "ideology" of science.

That Hull does not defend the status quo is testified by his interest on issues that a complete EET has to tackle (Hull 1988a chapter 14): how much competition is too much? Should we make science more efficient by making it more cooperative? How much conceptual heterogeneity is optimal in science? Can we reduce the number of scientists without reducing the pace of evolution? To give answers to these questions might be essential for any social epistemology. Work in this direction has been undertaken by many epistemologists, some of them not sharing EET's tenets.[135]

Despite the overtly critical overview of Hull's EET, I believe that he has contributed to the development of a more satisfactory and complete evolutionary epistemology. In three areas particularly, I see Hull's contribution to be innovative and defensible.

The first innovation of Hull's EET pertains to the explanatory strategy he provides, that is, its populationism. Hull's special emphasis on explanations in terms of group selection seems to me illuminating. Hypotheses about cultural group selection are an important aspect of any EET. I have rejected the explanatory completeness of Hull's approach, basically because invisible hand explanations are just one kind of selective explanation, and because nothing prevents one from coupling them with individualistic explanations of some kind (e.g. intentional, internalist etc.). I therefore do not see any incompatibility in principle between Hull's EET and other individualistic epistemologies. This amounts to claiming that explanatory pluralism is a better explanatory strategy than selective or internalist monism.

The second respect in which Hull's EET is original pertains to the hypothetisation of the causal influence of some social features of the scientific process. This implies that Hull's EET does not endorse an individualistic methodology. In section 5.1 I shall assess whether and how social factors can act as causal agents, especially by taking into consideration Campbell's notion of downward causation.

The third innovation of Hull's EET is given by the emphasis on the social aspect of the process of knowledge growth, and the correlated thesis of the adaptedness of the social

[135] cf. Kitcher (1993), A. Goldman "Epistemology and Cognition" etc.

structure of science. Hull's EET is a social epistemology based on the promissory note (or perhaps unsubstantiated optimistic tenet) that science is a functional system somehow adapted to achieving its institutional aims. I shall dedicate section 5.2.3 to assess whether normative recommendations can be derived from an adaptationist epistemology.

Chapter 5: the epistemological significance of EET

" Are you awed by the exquisite fit between organism and environment, and find in this fit a puzzle needing explanation? Does the power of visual perception to reveal the physical world seem so great as nearly to defy explanation? Do you marvel at the achievements of modern science, at the fit between scientific theories and the aspects of the world they purport to describe? Is this a puzzling achievement? Do you feel the need for an explanation as to how it could have come about?....I must concede that I do not have compelling evidence of fit in any logically entailing sense....in seeing the fit, and the puzzle, I do so only on the basis of presumptions which go beyond my capacity to verify or compellingly demonstrate to another person....For the three problems of fit that began this inquiry, and indeed for all problems of fit, there is available today only one explanatory paradigm: blind variation and selective retention....We believe in it because it describes a possible route, and because there are, at present at least, no rival explanatory theories."
Donald T. Campbell, "Unjustified Variation and Selective Retention in Scientific Discovery", 1974, pp. 139-142

In sections 1.2 and 1.3 I illustrated the challenges faced by naturalised epistemologies. In this chapter I will show how EET copes with the challenges. I will consider three basic issues: the normative value of EET (cf. section 5.2), its rebuttal of the relativistic challenge (cf. section 5.3), and its position on the issue of realism (cf. section 5.4).

So far I have argued that EEM models are largely irrelevant epistemologically if our aim is to understand the process of scientific change (cf. chapter 2). I have instead focused on different evolutionary models, which seem much more promising, that is, EET models (cf. sections 3.1. and 3.2). I have then analysed the two most famous and arguably best EET models so far proposed. In both cases I have tried to show that the evolutionary approach leads the two authors to embrace a sociological perspective whereby the scientific process is functionally analysed as if it were an adaptive and partly adapted system.

I shall now briefly characterise the nature of EET models by focusing on their basic and most important features.

First, the EET models all share a basic evolutionary commitment to the validity of the variation-selection model of knowledge acquisition. This explanatory paradigm can be applied as a psychological theory of hypothesis formation (cf. 3.4), or more generally as an empirical hypothesis concerning the nature of knowledge: any biological, psychological, cultural and social knowledge process is selective. For terminological reasons I shall refer to the "blind variation and selective retention" model by using the acronym BVST, despite "blind" not being an apt term to refer to the nature of intellectual variation (cf. again section 3.4). The reason is that this is how the model is generally referred to. I will focus on the epistemological nature of this commitment in subsection 5.2.2.

Second, genuine EET models seek an understanding of the scientific process not merely or mainly in analogical fashion (cf. section 4.7), but from a more general perspective that can be termed "general selection theory" (I will use the acronym GST - cf. section 3.3). This is because it is recognised that the analogical reasoning so preponderant in the early illustrations of evolutionary epistemology has no significant epistemological value (apart from providing sometimes useful heuristic hypotheses and metaphors). The nutshell is that science evolves because of some proximate social selective factors and not because it metaphorically evolves by natural selection (cf. section 5.2.3).

Third, EET becomes mainly a social epistemology. This is because of two main reasons. First, the selective mechanisms operative in science are social ones. This means that the traditional individualistic perspective to solve epistemological issues is abandoned in favour of the social one. The problem EET tries to solve is how communities of people, structured in whatever way, come to know. One implication of such a view is that some traditional epistemological problems (e.g. the problem of other minds) are not even treatable from EET's standpoint, because a solution to them is presupposed. Second, EET treats knowledge not abstractly but as carried by a vehicle (cf. section 1.3). In the case of scientific knowledge the vehicle of main epistemological interest is the community of scientists, however characterised.

Before treating the central issues concerned with the normative value of EET and its putative relativism we shall consider the nature of its methodology. Individualism as a reductionist thesis (both at the methodological and ontological level) has been and still is an unsubstantiated assumption concerning the proper way to conduct scientific research. In the next section I will show that, despite the fact that individualism is still the mark of respectability for scientific and epistemological research, EET does not endorse individualism.

5.1 Methodological populationism

In this section I shall show why I think that methodological individualism is not the correct kind of epistemological approach for an evolutionary and social epistemology as is EET. After illustrating the general nature of EET's alternative methodological approach (i.e. populationism), I will provide five arguments to substantiate my thesis.

Methodological populationism in biology and sociology

EET's theorists can choose between two alternative methodological approaches. The first is methodological individualism.[136] The other can be termed methodological socialism, collectivism, or populationism (the term I favour).

Prima facie, an epistemology, like EET, which is evolutionary and social, should be methodologically oriented as its sciences of reference are. In this sense, we could ask whether evolutionary biology and sociology are methodologically individualistic.[137] If they are not, then an a priori reason to prefer an individualistic epistemological framework for EET is rejected.

Many biologists have argued that the major novelty of Darwinism is the populational approach, originally envisaged as an alternative to essentialism (Mayr 1976 p. 293). In evolutionary biology the unit of evolution is the population. It is the population that evolves, that changes, not the individuals constituting the population (Lewontin 1982 p. 155). Natural

[136] Favoured by P. Kitcher (1993) ch.8.
[137] Populationism is, however, not only restricted to biology and sociology. What can be termed population-thinking is a general methodological approach with instantiations in many sciences (e.g. physics - theory of gases) whose

selection does not explain the existence of a certain trait in a certain organism, but only the frequency of that trait in a whole population. Natural selection is not concerned with the explanation of individual's characteristics because that is the province of ontogeny.[138] The basic tenets of population-thinking are that the population is the unit of biological organisation, and that the individual is relevant in the evolutionary process because it is individual differences (phenotypic variation is individually based) that are to be considered the engine of evolutionary change. In evolutionary theory, the description of evolutionary phenomena is by means of properties of populations, and evolutionary forces have effects on populations. The populational approach allows evolutionary theory to look for invariances at a higher level of biological organisation than the individual. Populationists are for this reason engaged in looking for invariant properties of populations in which there is systematic variation at the individual level. That is, population biology is based on the methodological assumption that individual properties can be understood in terms of population ones. Population properties in population biology are autonomous objects of inquiry.

This means that biology is not methodologically individualistic. However, there is no a priori link between populationism as a methodological and as an ontological thesis, in the sense that methodological populationists are not necessarily committed to the "reality" of population properties.[139] In fact, some biologists argue that population-thinking implies that only individuals are "real", while populational properties are not (Mayr 1976). Others reply that just because populational properties are measured statistically, this does not imply that they are only "abstract" artificial devices ontologically on a par with disparaged essentialist properties. This metaphysical implication of population-thinking should be rejected (Sober 1993 p. 165-8).

---

general justification is given by the belief that constructing theories in terms of populational properties is fruitful and pragmatically justified.

[138] Neander (1988) disagrees with this view.

[139] Mayr (1976) sees population-thinking in biology as the vindication of organismic selection, as an ontological hypothesis. In this sense individuals are the only "real" (i.e. causally efficacious) entities. The methodological implication of population-thinking as an ontological hypothesis would be, following Mayr's line of reasoning, reductionism: populational properties are to be reduced to organismic ones. Methodological reductionism is prima facie justified by the fact that organisms are easily recognisable entities. But they are still higher-level entities compared to cells and genes. Following methodological reductionism to the extreme would vindicate the gene-

In sociology, methodological individualism is, as Campbell puts it:

"...the dogma that all social group processes are to be explained by the laws of individual behavior - that groups and social organizations have no ontological reality - that where used, references to organizations, etc. are but convenient summaries of individual behavior."[140]

Campbell rejects methodological individualism both in his sociological and epistemological writings. For Campbell, EET is partly a social epistemology that requires an alternative methodological approach. This is because methodological individualism is an insufficient descriptive and explanatory strategy for epistemological investigation. One cannot do a good job of understanding science without population-thinking because:

"....the selective system of science is ultimately socially distributed in a way which any individualistic epistemology fails to describe adequately."[141]

I turn now to explain why I fully agree with Campbell, and why EET theorists should reject individualism.

Five "good reasons" for rejecting methodological individualism

Methodological individualism is a methodological thesis according to which the description and explanation of natural (comprehending social) phenomena (and also patterns or groups of phenomena) involving macro properties must be reduced to micro accounts (ultimately physical). Such a thesis has as a corollary the outcome that looking at phenomena from a macro perspective is not even pragmatically fruitful, since macro accounts neither disclose nor capture, for example, useful generalisations concerning patterns of events.

---

selectionist hypothesis (and in general genic reductionism). I doubt Mayr would be happy with this extreme implication of his own thesis.

[140] D. Campbell (1994) p. 23

[141] D. Campbell (1974a) p. 437

Methodological individualism is therefore coupled with an ontological series of theses. The reductionist thesis states that macro properties have no ontological reality, no causal efficacy. In this sense, macro accounts might be only more convenient, but do not capture an underlying reality apart from the ultimate micro (i.e. physical) reality. The supervenience thesis states that macro properties supervene on micro ones, in the sense that there cannot be a macro difference between two states of affairs without an underlying micro difference.

In epistemology methodological individualism states that the success of science can be explained purely in individualistic terms, by merely referring to the cognitive mental operations of individual scientists. Methodological individualism in epistemology states that scientific phenomena must be studied from an individualistic perspective because the explanation and description of science's success can be reduced to the successful selective decisions of individual scientists. Any reference to group properties or features of the social structure of science should be eschewed because explanatorily irrelevant, even as far as generalisations concerning patterns of scientific episodes are concerned. This means that the workings of the social system of science can be fully explained in individualistic terms. This also means that scientific selection acts only on individualistic properties.

To such strong theses concerning EET's descriptive and explanatory role some metaphysical theses are added. The macro properties (i.e. the non-individualistic properties to which EET's explanations refer) have no ontological reality and no causal efficacy. Furthermore, science's success completely supervenes on the individualistic mental properties of individual cognisers.

I believe that methodological individualism cannot be a sound epistemology for EET because the, I believe, genuine explanations EET provides do not respect methodologically individualistic standards. This is because of five reasons.

First, it is false that all episodes of success in science can be fully described or/and explained by merely referring to the mental operations of individual scientists that determined their successful selective decisions. This is because reference to properties of research groups or even larger communities is explanatorily relevant. This was shown in chapter 4 where I defended the thesis that Hull's explanation of why cladism succeeded was at least partially explanatory. This explanation showed that some scientific episodes can be explained by reference to properties of the research group, in the sense that group properties are "selectable" in the scientific process. Furthermore, such group properties cannot be reduced to individualistic ones, since the idiosyncratic behaviour of individual scientists cannot always be generalised upon. Cladism succeeded because cladists were theoretically tighter as a group, but the group property cannot be reduced to any peculiar individualistic "type" behaviour that all cladists performed without rendering the explanation less significant (cf. section 4.5).

Second, reference to social features and properties of the adaptive organization of science are explanatorily relevant. This means not only that particular social features of the scientific process have a causal effect in furthering science's goals, but also that their action cannot be reduced to individualistic properties, given that individuals do not mentally instantiate such features.

From EET's perspective, science is a teleological process also because its practice is governed by specific norms and by specific goals that exist independently of the goals of individual scientists (cf. sections 4.5 and 4.8). The social norms and science's institutional arrangements partially direct the sometimes idiosyncratic cognitive behaviour of individual scientists.

This was again shown in chapter 4. For instance, Hull points out that science's social system would be affected by giving up the convention at the heart of the credit system according to which all credit is given to the first publisher of some research. In this sense, the scientific process happens to be cumulative not only because scientists are driven by seeking credit and public recognition (i.e. not because they constantly remember that good research means celebration), but rather because such social convention exists, affecting scientists' behaviour.

From an ontological point of view the convention exists independently of any scientist, and also independently of any psychological instantiation. Rather, this social feature of the scientific process might causally affect, by feedback, scientists' behaviour.

Hull's explanation of why lying is rarer than stealing in science works in the same way, by identifying the social mechanism for the minimisation of lying (i.e. severe punishment) that prevents the scientific process from working ineffectively. There is not even reason to brainwash scientific trainees about the maladaptive function of lying, since mechanisms for the detection and punishment of this kind of behaviour are already operative in science.

Campbell's "anti-tribal" norms of science provide another example of structural properties of the scientific social system that are explanatorily relevant, and whose ontological reality is independent from any particular psychological instantiation. In this sense, social norms, as products of cultural group selection processes, have an ontological reality that is independent of the individuals who happen to endorse them, independent of their instantiation in the mind of particular individual cognisers. In this sense, social norms transcend the human capacity to mentally grasp them. More generally, social norms are an exemplification of what Campbell calls vicarious selectors, that is, of products of selective processes belonging to a higher level order of biological/cultural organisation which causally affect lower level selective processes (cf. section 3.3 and point 4 in this section).

Third, generalisations concerning patterns of scientific events couched in evolutionary and social terms can be genuinely explanatory. I have already pointed out that Hull's EET major explanatory value lies in describing and explaining what many historical scientific episodes have in common, rather than peculiarities and/or previously undisclosed causal factors influencing singular episodes.[142] This might be true for EET models in general. Examples of such generalisations are Campbell's hypotheses about the doubt-trust ratio and about the functional role of deviant behaviour in science (cf. subsection 5.1.3), claims concerning the

[142] Hull (1988a p. 388) contends that generalisations, hopefully in the form of laws of selection, can be captured, but only if we ascend the level of epistemological analysis. If the science of science is ever going to capture generalisations about the scientific process, then, Hull contends, we must look at a higher level of generality than the individual scientist, or even the research group, but rather at the level of kinds of groups and research programmes.

equilibrium relationship between variation and selection, plus all claims concerning the adaptive (or maladaptive) nature of the self-policing system in science and the general adaptedness (or maladaptedness) of the scientific process.

EET tries to capture invariances at the level of populations, that is, significant generalisations about the populational aspect of the scientific process. In order to fulfil this task, populational properties are used. These properties should, according to reductionists, be reduced to cognitive (wrong) and ultimately physical (right) ones. However, EET theorists should argue that, when we want to explain what patterns of scientific episodes have in common, then that explanation might refer to non-subvenient properties. The vocabulary of supervenient properties makes an irreducible contribution in this respect. This means that, even though the explanation of any singular historical scientific episode could be framed in terms of the subvenient properties (i.e. reductionism concerning token events), it does not follow that the explanation of what many episodes have in common can be couched completely without reference to the vocabulary of populational or supervenient properties (i.e. anti reductionism concerning type events – Sober 2000 section 3.5 & 1999b).

While we have seen that population-thinking seems to assume a radical holist aura in Hull's EET (cf. section 4.6),[143] I believe that EET should reject both methodological individualism and radical holism while embracing antireductionism, that is, the thesis according to which the explanation of the success of science must be couched both in populational and individualistic terms. In this sense, EET is a more complete epistemology than any individualistic one.

Fourth, there is no reason to believe that the social properties to which EET's explanations refer have no ontological reality and no causal efficacy. What Campbell and Hull show is that scientific practice is governed, among other things, by the existence of norms of scientific behaviour (e.g. institutional norms of science, ideological commitments, regulative ideas etc.), which regulate and direct scientists' behaviour. Furthermore, such norms have an independent ontological reality (cf. point 2).[144]

This view is generally rejected because it requires a "weird", according to the critics, notion of downward causation. However, downward causation is not equivalent to emergent causation. In this respect it must be said that physicalistic supervenience and the thesis of the causal completeness of physics are incompatible with the thesis of emergent causation, but do not entail that macro properties are causally inert.[145] This is exactly Campbell's (1974c) view. Campbell accepts the physicalistic version of the supervenience thesis, the thesis of the causal completeness of physics, but rejects the thesis of emergent causation. In brief, in Campbell's view the thesis of the causal efficacy of macro properties is not a thesis about emergence, but a thesis about (downward) causation.[146]

For Campbell, the thesis of downward causation is an epistemological thesis, a thesis regarding the proper and most fruitful description and explanation of events. He claims that reductionism must be extended as to encompass two further principles apart from the basic physicalistic assumptions (i.e. that higher level processes act in accordance to physical laws and require for their implementation physical mechanisms and processes, and that the explanation of the workings of higher level systems and processes requires reference to physical micromechanisms). The first is termed the "emergentist" principle, which states that biological and socio-cultural evolution encounter laws that operate as selective systems which are and will not be described by contemporary laws of physics or their future substitutes. The second additional principle is downward causation, which states that the laws of the higher level selective systems determine in part the distribution of lower level events and substances,

---

[143] Radical holism is the thesis according to which the causal efficacy of lower-level properties and processes is denied.
[144] By this I mean that such norms are independent of the minds of scientists, not of the more comprehensive physical environment (e.g. books).
[145] Cf. E. Sober (1999a), who rejects Kim's argument against non-reductive physicalism (as stated in "Supervenience and Mind" ch.s 14 and 17).
[146] Having said this, I am not sure whether evidence can adjudicate between emergent causation and physicalistic supervenience.

that is, that all processes at the lower level are restrained by and act in accordance to the laws of the higher level. Campbell adds that:

"'Downward causation' is perhaps an awkward term [for the last principle].....The 'causation' is downward only if substantial extents of time, covering several reproductive generations, are lumped as one instant for purposes of analysis. In the 'instantaneous' causation of the older philosophical analyses of physics, no such direction is present".[147]

From this passage it seems to me clear that Campbell accepts that downward causation is not the negation of physicalistic supervenience, which implies (or is equivalent to) a thesis about instantaneous or synchronic determination. The issue is whether downward causation is merely an epistemological thesis about how to explain and describe reality, or whether it is a metaphysical thesis about causality. Campbell thinks it is just an epistemological thesis, since all causation is direct physical causation. When Campbell claims that macro properties, however supervenient, act as restraints on the lower levels, and have for this reason causal efficacy, he is not saying that this causal role can be fulfilled independently of their supervenient basis (this would be emergent causation), but only that from an epistemological point of view purely physicalistic explanations cannot be complete. The epistemological version of the thesis is sufficient to give up methodological individualism and embrace an anti-reductionist attitude. As the biological explanation of why, in certain species of ants, soldiers have such large and specialised jaws that they cannot feed themselves must refer to the division of labour and facts about the general social organisation of the colony (Campbell 1974c p. 181), so explanations of why science is successful must refer to the influence of the vehicular contribution.

Turning back to EET, the validity of Campbell's epistemological thesis seems to me indubitable. That social norms and vicarious social selectors causally affect the outcome of the scientific process by restraining and directing the cognitive and selective operations of

[147] D. Campbell – "Downward Causation" p. 180

individual scientists seems to me obvious. Given downward causation, the practice of science would neither be describable nor explainable by reducing its operations to the actions of individual scientists, and the generalisations concerning science could not eschew reference to higher level entities such as groups, communities, institutional goals of science, and various social vicarious selectors. In this respect, EET does not commit the individualistic fallacy (i.e. the fallacy of considering ontologically "real" only the properties of individuals).

Fifth, the "cognitive" supervenience thesis is false, since success (e.g. genuine consensus formation) in science does not solely depend on the mental properties of individuals but also on the adaptedness of the social organisation of the scientific process, which can be guaranteed even without the psychological instantiation of any sound belief (given that certain social vicarious processes govern part of the system). This means that a higher level event (science's success) is not completely dependent, determined or necessitated (all characterisations of the supervenience relationship) by the cognitive and mental operations of scientists, and especially their "rational" selective choices. Science's success generally depends on additional events and states of affairs that are not cognitive (i.e. they do not have an individualistic cognitive instantiation). EET's explanations taken into account at points 1, 2 and 3 are of this sort.[148]

This does not mean that the physicalistic version of the supervenience thesis is not correct, namely that physical events and states of affairs at time t do not determine and necessitate the state of affairs that is described by using macro (i.e. non-cognitive, that is, for instance, social) properties. Giving up the cognitive version of the supervenience thesis has no implications for

[148] The supervenince claim states that the success of science can be reduced to the causal action of the rational decisions of agents, while my claim is that science's success is "also" a function of the social structure of the scientific process. This means that social and populational properties have causal powers that transcend those of epistemic agents. How is this possible? We might appeal to a counterfactual notion of causation: the social properties are such that, had they been different or absent, the effects would have been different. I believe that such a notion is not needed. This is not because I have a problem with not easily testable counterfactual claims, but because there is actual empirical evidence to justify my claim: the history of the French Academy in the seventeenth century, when the winners-take-it-all convention was briefly given up (section 4.3), shows that progress (e.g. cumulativity) was impaired because scientists kept results secret. Of course, it might be argued, this is because the convention was mentally instantiated in some form in the minds of scientists. If this is true, then some form of the cognitive supervenience thesis is correct, but this fact (upon which I am of a double-mind) does not show that methodological

the validity of its physicalistic version. To put the same point in another way, if populationism is a correct methodological thesis we have an argument against cognitive supervenience. This is because EET's social, selective and populational explanations of the growth of knowledge might provide better explanations than purely individualistic ones. These are the reasons:

1) non-cognitive properties are somehow causally efficacious (downward causation); this is point 4;

2) EET provides generalisations concerning patterns of events and the features that various episodes in science have in common that cannot be captured from an individualistic perspective (point 3).[149]

I believe that any epistemology that endorses the cognitive version of the supervenient thesis (e.g. Ruse's EEM, cf. chapter 2) is fallacious in this respect, since it suffers from an individualistic bias. If social selective processes are causally operative in science, and if the workings of these processes are not cognitively supervenient on the individual behaviour of scientists, then part of the explanation in such cases relies on social selective processes. For this reason social epistemology is a fundamental part of a complete epistemology of science.

In this section I have tried to show how EET offers an alternative to individualistic epistemologies. In the next section I shall focus on its normative agenda.

5.2 The normative value of EET

In this section I wish to characterise the ways in which EET models are normatively significant. In chapter 1 we saw that naturalists disengage from the traditional normative

individualism is correct, while it certainly is not evidence in favour of the rationalist version of the supervenience thesis.

[149] If populationism is a correct ontological thesis, then cognitive supervenience fails for the additional reason that not all properties that exist and are "real" are cognitive, with perhaps the consequence that higher-level properties cannot be reduced to physical ones (emergent causation). Only if emergent causation exists physicalistic superveninence fails.

agenda and thus give up the programme of epistemic justification as traditionally conceived. I believe that genuine naturalists should embrace what Kitcher labelled the "meliorative project" (cf. chapter 1), that is, the attempt to improve our cognitive performance in the actual world. In this sense, the central problem of epistemic naturalism becomes the maximisation of epistemic utility for cognitively limited creatures in the actual world (Kitcher 1992 p.24), rather than some unachievable, obscure and obsolete epistemic quest (e.g. certainty). Furthermore, by rejecting the legitimacy of basing epistemology on a priori knowledge, what is left to naturalists is hypothetical normativity, that is, a theory of science that is based on empirical and hence never fully proven assumptions about the world to be known.

EET is clearly committed to the relinquishment of the traditional normative project. The alternative endorsed is, as already seen, hypothetical naturalism. EET's normative claims are conditional on various naturalistic assumptions regarding the cognitive capacities of human beings, the theory of the functional analysis of social systems, and, more generally, the validity of the variation-selection model of the emergence of adaptation. EET theorists think that naturalistic justification is a matter concerned with the function and aim-conduciveness of the behaviour prescribed by the norm. Recommendations can be derived regarding ways to better pursue the aims of science, by both affecting scientists' behaviour and by altering the epistemic conduciveness of the scientific process. For these reasons EET certainly has the potential to fulfil the "meliorative" role central to the epistemological endeavour. The thesis I shall defend in this section is that EET has many different kinds of normative insights to offer.

The nature of EET's normative project can be thus outlined. From the analyses proposed in chapter 3 and 4 we have seen that EET studies knowledge not abstractly by eschewing reference to its carrier or vehicle (e.g. the epistemic agent), but rather by considering the vehicle's features. From EET's perspective both logical and individualistic analyses of scientific change are incomplete, while the main carrier of epistemological importance is social (e.g. the community of scientists). We have also seen that EET has two epistemological souls. The first is the evolutionary, the second is the social.

In the first case, EET tries to find out about the scientific process as an evolutionary process, and tries to come up with insights concerning the most adapted and effective organisation of the scientific process. In order to do so, EET relies, like other naturalistic approaches, on some previous knowledge in the form of working hypotheses (or methodological assumptions). The basic and most general one is the thesis of the universality of selection, stating that a blind variation and selective retention process (BVSR) is fundamental to all inductive achievements, to all increases in knowledge and to all increases of fit between representation and represented (cf. section 3.3). One way to interpret this claim normatively will be seen in section 5.2.2.

EET as a social epistemology puts special emphasis on the analysis of the social structure of science (e.g. its self-policing structure - Hull esp. cf. section 4.2 - and its ideology and social norms - Campbell cf. 4.8). EET theorists will be particularly concerned with the social validation of scientific knowledge and with the effects of the social organisation of science on the growth of knowledge. EET has the normative role of identifying the adaptive and epistemically conducive features of the social system of science, a normative task completely lacked by epistemologies that are not social. The normative recommendations EET as a social theory of science tries to capture concern the individuation of the institutional arrangements and social norms designed to maximise the selective role of nature (cf. section 5.2.3).

Evolutionary and social analyses merge naturally because the social system of science can be studied in terms of functional organisation and adaptation. Before moving to assess EET's project in detail we need to focus on one of the main alleged obstacles in the path to achieving genuine normativity: the naturalistic fallacy.

5.2.1 - Validation of the norms of science and the naturalistic fallacy

Hume, according to whom normative statements cannot be "deduced" from descriptive or empirical ones, first highlighted the naturalistic fallacy.[150] It is generally agreed that

[150] More precisely, I am here referring to the so-called "is-ought problem" and not to Moore's fallacy.

deducibility is too strict a relationship between the two kinds of statements. What is sufficient is the existence of an inferential relationship between the two. The fallacy potentially affects all naturalistic epistemologies that base their analysis on factual claims.

In this section I will first clarify how the naturalistic fallacy could affect naturalistic and EET models. In particular, two versions of the fallacy are potentially threatening. My opinion is that they are not damaging at all and that EET does not commit the naturalistic fallacy. Thus, its normative project cannot be criticised for this reason. However, in the course of the analysis we will return to consider a more serious challenge of naturalism already considered in section 1.2.1.

The weak inferential version of the naturalistic fallacy states that is-premises cannot provide non-deductive evidence for ought-conclusions. Let us assess this claim by comparing two naturalistic solutions.

One of the most articulated naturalistic attempts to solve the fallacy has been proposed by Laudan (1987a and 1996). His account, as EET's, is based on a few typically naturalistic premises (cf. section 1.1). First, naturalists believe that epistemology does not require categorical but only conditional norms (Giere 2001 pp. 57-60). Naturalists think of epistemic norms as hypothetical imperatives based on factual and empirically assessable statements. Secondly, conditional or hypothetical norms can be derived from descriptive statements and can be justified naturalistically, that is, empirically. As I argued in section 1.1, I believe that naturalists should reduce epistemic to instrumental value so that epistemic "oughts" become identical with descriptive facts about instrumentally appropriate behaviours relative to epistemic ends. Laudan remains agnostic about the idea of reduction but prefers to see the fact-value link as a dependence relationship.

Laudan's strategy to justify norms is as follows. Suppose that our epistemic aim is to achieve goal A. The normative claim will be relativised to the efficacy of different methods (for simplicity let us consider only two, x and y) in achieving A. The conditional and hypothetical norm:

a) If your aim is A you "ought to do" x.

is dependent on the empirical claim:

b) Doing x is better than y in order to achieve A.

Laudan links normative and empirical statements, where the latter provide evidence in favour of the former. As a consequence, the facts of the world have epistemological implications. It seems to me that Laudan does not commit the naturalistic fallacy because he clearly shows how is-premises provide non-deductive evidence for ought-conclusions.

EET adopts the same approach. Its justificatory strategy starts from considerations of this kind: given that goal A has been chosen as our epistemic aim, can EET tell us what means will further the achievement of such a goal? The structure of the argument is the following. First premise: A (e.g. empirical adequacy) is the goal of science. Second premise: social feature or behaviour x is more epistemically conducive than an alternative y. From these premises it follows that x is preferable to y, and that we should behave so as to maximise the influence of x and minimise that of y. To give an example:

(A) If your aim is to achieve empirically adequate hypotheses, then you ought to respect the anti-tribal norms of science.

(B) Conforming to such norms is better than behaving in accordance to the authority of the Bible in order to achieve empirically adequate hypotheses.

As in Laudan's case, it seems to me clear enough that EET also does not commit the fallacy. I conclude that, if the naturalistic fallacy simply states that is-premises cannot provide non-deductive evidence for ought-conclusions, it is not a significant intellectual pitfall. This is because it is obvious that the facts of the world have epistemological (and ethical) implications.[151] Let us see if a stronger interpretation of the fallacy can seriously challenge the naturalistic approach.

According to this stronger interpretation, the fallacy states that is-premises cannot provide non-deductive evidence for ought-conclusions *unless normative auxiliary assumptions are added to link them.*[152] To give an ethical example: from the descriptive claim that smoking causes health problems I cannot conclude that smoking in public places is immoral unless I add the normative claim "Causing health problems to people by smoking in public places is immoral".

But where does the problem lie in the epistemological case? The structure of the naturalistic argument is the following:

(Factual premise) Doing x leads to the achievement of A better than alternative y

(Normative assumption) Achieving A is cognitively fruitful

(Conclusion) Doing x is cognitively fruitful

My suspicion is that there are two problems with this argument. The first concerns the factual premise, which presupposes a theory of comparative testing. The second problem regards the justification of the normative assumption, which presupposes an axiological

[151] If I know that chicken is unsafe to eat, then I think this is a very good reason to try to convince people that they ought to avoid it.

[152] E. Sober "Prospects for an evolutionary ethics" p.109. Even if Sober's argument regards ethics, I believe it can be extended to epistemology, given the structural similarity of the naturalistic fallacy in the two cases.

theory. Let us first consider how Laudan solves the two problems highlighted above. Then we will move back to EET.

In Laudan's case we can ask: isn't claim (b) sufficient to derive claim (a)? The answer is, partially, yes. The problem is that claim (b) is a comparative statement implicitly based on another one, namely:

c) claims of type (b) are justified if and only if produced by a good method of investigation

where the good method of investigation is the basic rule of enumerative induction R1 (cf. section 1.1). Thus, any inductive method used to assess claims of type (b) that can be reduced to R1 will be "justified". By means of this justified method we start gathering historical evidence as to the aim-conduciveness of methods x and y and compare them. Laudan's solution to the first problem is that confirmed empirical statements of type (b) provide a basis to derive and justify claims of type (a). The problem becomes that of justifying reliance on R1. Regarding the axiological problem, Laudan thinks that problem-solving effectiveness is the aim of science.

EET's justification procedure is structurally similar to Laudan's. Premise (B) in EET's argument is justified via the following statement:

C) claims of type (B) are justified if and only if produced by a reliable method

where the reliable method is eliminative induction. With this method we collect historical evidence concerning the relative aim-conduciveness of different social practices and by elimination we retain the best ones. In this way ought conclusions of kind (A) are derived from is premises while the latter assume a justificatory function. Of course, the methodological problem is only apparently solved because what EET theorists need to show is that eliminative induction is a reliable method. I shall focus on this issue in the following sub-section (5.2.2). The axiological problem is justified by appealing to a certain definition of epistemic conduciveness that is compatible with our capabilities and epistemic limitations. This definition might in turn be justified by appealing to facts about our culture and socialisation (e.g. that we convened to choose aim A), or perhaps even in terms of moral principles (e.g. the idea of truth as a duty). This process is empirical, partly regarding the cultural advantages that a community of agents reaped by choosing one epistemic aim rather than another (this is a cultural group selection process – cf. Sober and Wilson 1998 chapters 4 and 5). I personally believe that EET can aspire to define the cognitive good in instrumentalist terms, but opinions diverge. I shall return to this issue in section 5.4 where I shall consider whether EET should be committed to axiological realism.

I conclude that if the fallacy is interpreted as stating that is-premises cannot provide non-deductive evidence for ought-conclusions *unless normative auxiliary assumptions are added to link them*, then, again, EET does not commit the fallacy. This is because EET adds normative auxiliary assumptions to link factual premises and normative conclusions (and to ultimately justify the latter). The real problem is rather the circularity of EET's approach: how do EET theorists justify the reliance on the methodological and axiological auxiliary normative assumptions they use to link facts and values? Part of the answer will be supplied in sections 5.2.2 and 5.4 respectively.

To summarise, EET theorists believe that the only way to validate epistemic standards is by showing that norms are epistemically conducive (i.e. instrumental in leading to the achievement of the chosen aims of enquiry). This is because EET theorists, like fellow naturalists, think that norms can only be tentatively validated via empirical means and that no privileged a priori path can provide an alternative credible justificatory procedure. The only alternative to a priori justification remains that of deriving ought from is statements. In EET's case, the validation process is based on many working hypotheses concerning the nature of the scientific process. From the presumptive knowledge given by these hypotheses, the EET theorist tries to understand what function a particular kind of behaviour or a particular social feature of the scientific process has, and then determines whether such behaviour or institutional arrangement will be epistemically conducive. EET thus relies on a theory of epistemic conduciveness. The important problem concerns the nature of such theory. It is here that naturalists and anti-naturalists mainly disagree. In fact, traditionalists believe that the basic judgements about the epistemic conduciveness of, for instance, methodological rules, are both a priori and detached from the science they are designed to validate. Traditionalists favour the a priori path to justification because the process of naturalistic justification is circular. So the real problem is the circularity of the EET approach, not the naturalistic fallacy. We thus return to the issue treated in section 1.2.1: can EET solve the circle by naturalistic means?

Before answering this crucial question we need to be clear about the nature of the circle affecting EET. The following is one way to illustrate it. EET aims to discover the selective mechanisms that render science aim-conducive. In order to pursue this investigation EET looks at "successful" scientific practices (e.g. biology as practiced by the community of taxonomists, as Hull 1988b chose) and tries to understand which selective mechanisms were operative and causally efficient in those cases. But the term "successful" is already a normative term. So the question becomes: how do EET theorists decide whether the practices they study have been successful? EET theorists are forced to reply that the practices they study are considered successful because they show the causal influence of the selective mechanisms they consider being aim-conducive, where this reply is circular. It is circular because, in order to justify the analysis of particular scientific practices, EET theorists need to rely on some previous normative commitments (i.e. the normative auxiliary assumptions I refer to that include a theory of what constitutes reliable knowledge that led them to choose a particular practice to analyse, a theory of induction by means of which the evidence concerning the aim-conduciveness of the selective mechanism is gathered, and a theory of what constitutes epistemic value).

I call this set of normative assumptions "theory" because I believe that this knowledge is theoretical and part of the scientific epistemology EET tries to build. These assumptions are conjectures and they are therefore tentative and open to revision. EET's justificatory procedure is the only viable naturalistic solution of the circle: tentatively trust most of our beliefs about what successful science amounts to in order to assess our claims about the aim-conduciveness of the extrapolated norms. At the same time confidence in the norms hopefully increases if it can identify new successful practices or new instances of successful science. The constant feedback between the social mechanisms at the basis of successful science EET identifies and the evidence cited in favour of their aim-conduciveness exemplifies the more general process of reflective equilibrium between types of inference and successful instances characteristic of any naturalistic approach. The circularity exists but its pitfalls can be bypassed through the process of revision. I am not claiming that the circle is virtuous, but that the process of reflective equilibrium makes the process of revision possible. My optimism is confirmed by the fact that our standards of evaluation have changed and improved with time. However, some traditionalists argue that any attempt to exploit the reflective equilibrium strategy I have delineated must endorse some minimal normative requirements that are not open to revision. A healthy minimal a priorism, they contend, has no alternative and should be embraced by all naturalists (Worrall 1999). I partly agree with this. I fully agree that the revision process to which epistemic norms are subjected is based on some standards that are more central and less negotiable. But I am not prone to give up the concept of revisability in principle, even for the most central standards, even, for example, modus ponens. For this reason I believe that the minimal normative theory is part of science. More particularly, part of its most theoretical part. In section 1.2.1 I argued that ultimately the difference between the dogmatist and the "sensible" naturalist is pragmatical, that it concerns intellectual honesty. Naturalists have only to restate that humans have no alternative to the circular justification of knowledge, at least because a priori dogmatism is not a viable alternative. While dogmatists should answer this question: if dogmatic a priorism and circular naturalism are both intellectual vices, what are the reasons to prefer the first? My reason to prefer the latter is the optimism I referred to above: the possibility to correct and improve our epistemology.

I now pass to consider in more detail the nature of the theory of induction at the heart of EET models.

5.2.2 - The logic of the variation-selection model

The first respect in which EET models are normatively significant should be dear to traditional epistemologists concerned solely with the characterisation of the logic of science. In this section I will argue that there is a sense in which the blind variation and selective retention model of knowledge acquisition (BVSR) at the heart of EET models approximates this logic.

EET models are generally based on the thesis that the best possible strategy of investigating the unknown is via the "blind" origination of variants and their subsequent selective elimination and/or retention (cf. sections 3.3, 3.4, 4.5).[153] This view presupposes that blind variation and selective retention processes are somehow universal (cf. D. Campbell 1974b, R. Dawkins 1980, G. Cziko 1995, D. Dennett 1996). One way to interpret the universality claim is that BVSR processes are operative in all cases of knowledge acquisition and emergence of adaptation, where this thesis is meant extensively and is not limited to particular areas of investigation and/or parts of nature. An articulation of this conception is Dawkins' idea of universal Darwinism (i.e. the thesis that natural selection is a basic and primitive force in the evolution of the universe). Another articulation of the universality claim is Dennett's idea of natural selection as an algorithm, where the algorithm can be physically realised in a variety of ways. Today BVSR theorists accept that the universality thesis has some limitations and that

[153] I have already explained that "blind" is not the best term to use in this context (cf. section 3.4), because the production of biological, behavioural and especially intellectual variation is not purely random, but biased. The term blind does neither suit biological phenomena (only things that can see can be blind; genes do not see but mutate nonetheless), nor cultural ones. Campbell proposed "unjustified" as an alternative, but this term does not fit biological phenomena (why mutation should be justified?). I will use the term "blind" nonetheless for simplicity, since the model is generally labelled in this way in the literature.

non-selective processes (e.g. functional organisation) are operative in many cases. However, the universality thesis remains the default working hypothesis wherever adaptation is encountered (Cziko 2001).

The best formulation of the conception is due to Campbell, who argued that in order to have knowledge we need at least three conditions:

a) processes for the formation of variants, that is, a heterogeneity of alterations in the form of existing objects of whatever nature (e.g. ideas or memes, hypotheses, behaviours, genes);
b) processes for the selection via systematic eliminations of the variants;
c) processes for the preservation and retention of the selected variants.

Campbell famously argued that the process of BVSR provides the only available explanation of creative processes at all levels of biological and cultural organisation. An application of this view was illustrated in section 3.4, where I treated EET's psychological theory of hypothesis formation. BVSR, I argued in that section, provides a naturalistic approximation of the ideal "logic of discovery" empiricists had been looking for before fallibilist ideas emerged (Laudan 1980). More importantly for this section, I believe that BVSR also characterises the "logic" of justification. There is nothing disconcerting with this bold claim if the nature of BVSR processes is understood properly. In fact, BVSR processes have a variational and a selective aspect, where these aspects can be seen as analogous to the contexts of discovery and justification respectively. It is for this reason that the BVSR model at least *prima facie* provides a complete characterisation of the "logic of knowing".[154] My claim is even less surprising if we consider that there is a deep analogy between BVSR model and inductive eliminativism, and that the latter has been considered to provide a sound characterisation of scientific method by both naturalist (e.g. Kitcher 1993) and non-naturalist philosophers of science alike (most famously Popper 1972). In this section I shall argue that the general validity of the BVSR model can be defended. In order to do so I first need to clarify the nature of the relationships between BVSR and inductive eliminativism (cf. section 1.2).

The first thing to note is that BVSR and eliminative induction are not equivalent processes if the latter is intended as a "deductively" valid procedure. In fact, deductively valid eliminative induction means producing all possible hypotheses and eliminating all but one, which must be the true one. The BVSR processes operative in science cannot achieve such a logical rigour because they are natural processes. In particular, the BVSR processes at work in science that EET postulates cannot produce all possible variants, while the mechanisms of selection and preservation of the variants are not immune to error.

However, I have already tried to show in section 3.4 that the BVSR model offers the best explanation of creative processes. In that section I argued that, ultimately, all episodes of creative thinking can be reduced to the basic methods of production of variation postulated by EET, namely idea recombination and error in the replication of ideas.[155] If the hypothesis I defended in section 3.4 is correct, then it becomes obvious that in scientific practice there is no better alternative to the BVSR model in order to explain the emergence of hypotheses. This also means that the hypothesis presented in that section has a clear prescriptive element (i.e. the best method to introduce variation in a conservative process like science is via BVSR mechanisms – an instance of this recommendation will be presented in the next section). Of course, with this I do not mean to claim that via BVSR all possible variants will be generated.[156] The ideal of completeness has to be given up. Still, to restate the point of section 3.4, the best strategy (i.e. the one that most approximates the optimum, that is, the creation of all possible alternatives) to generate new hypotheses is via BVSR (this is the sense in which the model is the best approximation to the ideal logic of discovery).

The point of interest as far as EET is concerned is that, given that science is a natural process, the deductive interpretation of eliminative induction is, from a naturalistic point of view, partly irrelevant. This is because naturalists have, as their aim, the improvement of cognitive

[154] Of course, the term "logic" is confusing given that EET is a naturalistic approach. It would be better to say that BVSR provides a "science of knowing" (including the psychology and sociology of research).

[155] Note that at the variational level methodological individualism is valid. Note also that the idea of error in the replication of ideas is somehow captured by the notion of blindness.

[156] For a more optimistic view see Dawkins (1999 pp. 43-50), who claims that given sufficient time biological mutation will produce a copy of the optimal genotype for natural selection to pick up.

performance in the actual, and pursuing an unachievable cognitive goal does certainly not make things easier. However, deductive eliminativism remains an epistemological ideal.

A second clarification required concerns the way in which I am characterising BVSR processes. In this section I am mainly referring to the logical skeleton of the model, not to the model as it happens to be realised in scientific practice, which will be the province of the next section. From a logical point of view, the model is like an empty husk ready to be filled with mechanisms for producing variation, as well as for the selection and preservation of variants. Even though the mechanisms for producing variation are ultimately reducible to the two basic mechanisms of recombination and replicative error, in the next section I will show more complex social ways to feed the scientific process with intellectual variants. In the next section I shall also focus on the selective level. There is of course a big difference between filling the BVSR husk with modus tollens as the only selective mechanism and the EET social selective approach.

A third clarification required concerns the link between the BVSR model and eliminative induction. The theses are distinct but intimately connected. I consider the model as a stronger thesis in scope. To argue that scientific method is a BVSR process (the claim made by Popper and EET theorists) is not equivalent to arguing that it can be characterised in terms of inductive elimination (the claim sustained by Kitcher). The basic difference is that the latter characterisation pertains merely to the selective aspect of the scientific process, while the BVSR comprehends the variational and preservation aspect. In this section my aim is to defend the validity of the BVSR model. This implicitly involves defending the validity of eliminative induction.

A final clarification concerns the nature of the claim that scientific method can be characterised via the BVSR model. The claim can be understood descriptively or prescriptively. EET's stance is that both interpretations are appropriate. This is also how Popper interprets the model, and how Kitcher interprets the thesis that the selective process of science is eliminative induction. The interesting question is whether the prescriptive role of the model can be vindicated. In order to evaluate this issue I will focus on Popper's, Kitcher's and Campbell's attempts to "justify" the general validity of the BVSR model or of its more circumscribed selective element.

Popper famously claimed that trial and error (his name for BVSR) characterises the logic of the growth of knowledge (Popper 1972). Popper also famously claimed that such a method is not inductive. While the latter claim has been strongly criticised, the first has been treated seriously.

Popper famously asserted that he was not interested on the topic of the generation of hypotheses, showing a lack of interest on the variational aspect of the BVSR model. As a matter of fact, I consider his proposals to formulate bold and non-ad-hoc conjectures as recommendations pertaining to the variational side of the model, as mechanisms for the production of plausible hypotheses ready to be selected.

Concerning the selective level of BVSR, it is usually contended that Popper did not succeed in providing any sound "argument" for showing that trial and error is the best method of knowledge acquisition. An "argument" would amount, for instance, to show that trial and error leads to the selection of conjectures that have a higher degree of verisimilitude. Popper did not show this. What remains of Popper's argument is just the skeleton of an acknowledgement that such method can neither be strictly "justified" nor said to be "rational", where by "strictly" I mean that the reliability of the method can be logically demonstrated. Furthermore, Popper was not very precise in the articulation of the selective level. Sometimes he referred to *modus tollens* as providing the only sound rule of methodology, while other times he added that severe testing was necessary, without however fully explain in what the severity consisted. This said, I believe that a kernel of truth can be salvaged from Popper's attempt to make a case in favour of the thesis that, if any method is rational, that must be trial and error (of course with the provisos that all plausible alternative hypotheses are generated and evaluated, and that Duhem-style underdetermination is circumvented).

Kitcher's is not an attempt to justify the BVSR model, but a more circumscribed attempt to justify its selective side. Kitcher has recently argued that the familiar problems of induction can be solved by endorsing the eliminativist perspective. His argument is not aimed, like Popper's, at establishing the rationality of the method in principle, but is naturalistic. What Kitcher (1993 p. 241) proposes is a variant of the argument from natural selection considered in section 2.4. The argument tries to show that the human tendency to reason in eliminative terms is an adaptation. This evolutionary hypothesis about the origin of our eliminative inductive propensity provides a limited argument in favour of its reliability (how could we have survived so long as a species if our inductive habits are so fallacious?). However, despite the fact that the argument from natural selection provides a limited answer to the sceptic, it fails to justify science, for the reasons highlighted in chapter 2.

Campbell (1974a p. 436) spoke of the universality of BVSR as the "dogma" of his approach. He claimed that it is an "analytic" truth that in order to extend knowledge beyond what is already known BVSR processes are needed. If the BVSR model is truly endorsed dogmatically, then a further justification would not be needed. I have been arguing that this would not amount to a fully naturalistic option (cf. section 5.2.1). As a matter of fact, Campbell's view is "dogmatic" in name only, since the universality thesis is a falsifiable hypothesis. For instance, if the nature of intellectual variation is truly "Lamarckian" as contended by critics of the model (an issue considered in section 3.4), the proposal collapses.
The attempts to vindicate the universality and rationality of BVSR processes were given special emphasis in Campbell's early writings (until c 1988), while they vanished in later writings (cf. Campbell 1997).[157] Campbell gradually came to realise that any attempt to justify in general terms the rationality of BVSR processes is both impossible (it would amount to solve naturalistically Hume's problem) and, more importantly, irrelevant

[157] See Hull 2001e

epistemologically. This is because, Campbell thought, even though it is shown that the scientific process consists of a hierarchy of BVSR processes, this *per se* does not provide any justification of scientific knowledge. Campbell identified three reasons for this pessimistic conclusion.

First, even though the BVSR processes operative in science were shown to be adapted, this would not provide evidence in favour of their truth-conduciveness because adaptation only guarantees instrumental utility. For a scientific realist like Campbell this was a major setback (on this issue cf. section 5.4).

Secondly, not only scientific beliefs are products of selection processes, but also, for instance, religious ones. So, the reason why science is a superior epistemic practice depends on the nature of the selective processes involved, not on their being selective (on this issue see next section and section 5.3).

Thirdly, many selection processes operative in the scientific process are, as Campbell put it, antagonistic to competence of reference (i.e. they prevent nature to contribute to the co-selection of scientific beliefs – on this issue cf. next section)).

All these reasons identify one concern and a switch of emphasis. The concern regards the generality of the BVSR model, which, according to the later Campbell, renders it epistemologically superfluous. The switch regards the strong emphasis put on the social aspect of the scientific process. The conclusion Campbell (1997) reached is that selection theory on its own does not contribute to the validation of scientific knowledge, but that such validity rather stems from the social structure of science, and in particular from those BVSR mechanisms that allow competence of reference or referent co-selection. The new epistemological problem concerns the proximate causal role of the social structure of science, not the ultimate fact that the processes constituting such structure are BVSR ones. Campbell became thus engaged in developing a "sociology of scientific validity" (cf. subsection 5.2.3), instead of generally justifying eliminative induction or the validity of the BVSR model. But his sociological switch also amounts to an implicit abandoning of the evolutionary agenda.

Despite Campbell's pessimism, I believe that we can still save an important moral from Kitcher's and Popper's attempts to establish the epistemic legitimacy of some form of inductive elimination. In this sense my version of EET departs from Campbell's (and Hull's) because I believe that the BVSR model provides some epistemological warrant. I now move to explain what I mean.

BVSR processes have one major limit: they cannot be shown to yield the optimal epistemic outcome. This was also the moral of chapter 2, where I showed that natural selection (the epitome of BVSR) is not an optimising process. In order to achieve the optimum some variational and selective conditions must be satisfied. Focusing on the variation level, what we would need is a process generating all possible variants. In practice this ideal cannot be achieved. Also, in practice what counts are "plausible" variants. The problem is that if you generate all possible variants you know *a priori* that also all plausible variants have been formulated. The problem to be solved in scientific practice does not so much regard the quantitative aspect of the generative process, but its quality. What scientists should aim at is not the production maximisation of the number of alternative hypotheses compatible with the evidence, but rather the maximisation of the number of plausible hypotheses. Is there any chance to achieve this aim in practice? Despite a certain amount of motivated scepticism on the part of philosophers of science (mainly due to a strictly deductive interpretation of the idea of eliminative induction), scientists have many times stressed that the underdetermination thesis (i.e. the idea that indefinitely many hypotheses are always compatible with the available evidence) is a product of the under-representation of scientific practice: in practice, to produce one serious alternative hypothesis to the accepted one is already a success.[158] This suggests that even though logically the eliminativist idea is deficient, it may still work in practice.

Moving to the selective level, we can see that more problems potentially affect the BVSR model. Starting with the most general, it is argued that the commitment to eliminativism is not

[158] Cf. Kitcher 1993 pp. 244 ff

enough to "secure" the objectivity of scientific knowledge (Rosenberg 1996), and that eliminativism does not lead "demonstrably" to scientific progress (Howson 2000 chapter 5). I believe that even these critical points are generally outcomes of a purely logical interpretation of the idea of eliminative induction. A different opinion is expressed by some scientists (especially biologists, e.g. S.J. Gould), who suggest that, even though conclusive falsification is doomed from a logical point, it is not in practice.[159] I also acknowledge that the BVSR model faces many other challenges (e.g. the error-proneness of the background constraints used for the generation of hypotheses which would lead to the ramification of error during the selective process) that cannot be dismissed as pseudo-problems.

In a nutshell, in order to produce adaptations, BVSR processes require innovative variants and ingenuous selective procedures. At the scientific level this means that there is no alternative to human creativity and ingenuity in coming up with creative conjectures and creative testing procedures. But even an insanely optimistic belief about the epistemic capacities of humans is not sufficient to justify the thesis that BVSR is an optimal strategy to investigate nature.

I think that all these critical points serve to highlight the oversimplified nature of the BVSR model when applied to science, but do not affect its important residue of truth and do not undercut its highly relevant epistemological role. Thus, I share Campbell's pessimism in one respect: there does not seem to be any way to show that the BVSR processes operative in science are optimising, where this means that they will demonstrably lead us towards the correct picture of reality.

The argument so far has been that the BVSR model has many limits and that epistemic optimality (in the sense of eliminating without error all possible rivals) cannot be achieved. Given such a negative outlook, how can I vindicate the goal of the section and extrapolate a positive conclusion in order to defend the BVSR model? EET is a naturalistic approach committed to the pursuit of the meliorative project. Unrealisable cognitive aims are banned from its axiology because they are irrelevant if our task is to improve cognitive performance in

[159] Note that falsificationism is still the biologists' favoured methodology (cf. especially Hull 1988a chapter 6).

the actual. However, even though optimality is unachievable, we can still aim at it. As a matter of fact, I think optimality plays an important role in the EET framework as it plays an important role in biology.[160] It can serve as an epistemic ideal, as a regulative idea, as a goal worth approximating. In this sense, it provides a comparative measure to judge the relative adequacy of different cognitive approaches. My claim is that the BVSR model characterised in the skeletal terms provided in this section is, when appropriately filled in, the best available approximation to the optimal strategy to investigate nature. This means that if knowledge is produced by BVSR processes then this fact *per se* provides some kind of epistemological warrant for the reason that the beliefs so produced will be generated by a process that most approximates the optimal strategy to achieve the attainment of knowledge.

Note that, on the one hand, EET's BVSR model approximates what more traditional epistemologies recommend, and, on the other, the model provides a naturalistic interpretation of the validity of eliminative induction. In the first sense inductive elimination significantly characterises the basic principle of methodology according to which hypotheses should be tested against all possible rivals (e.g. Worrall 1999), further vindicating the thesis of the complementarity between EET's stance and that of some traditional epistemologies. In the second sense, the BVSR model provides an alternative to inductive foundationalism (which is incompatible with naturalism) in order to justify EET's commitment to eliminativism, the reason being that the selective process EET postulates aims at approximating the deductively valid version of eliminative induction.

Of course, the BVSR model must be properly characterised in order to show that it *actually* is the most reliable process of knowledge acquisition and that it is the one that most approximates the optimal model. The ideal is to have optimal mechanisms for the production of intellectual variants, optimal selective mechanisms and optimal preservation mechanisms. As I shall try to show in the next section EET aims at specifying the nature of some of these

[160] It is sufficient to think of the role played by optimality models to test adaptative hypotheses and in the field of

mechanisms. I think that EET partly succeeds in this respect, since it shows that the social BVSR processes operative in science contribute to improve our picture of reality. This means that, given that EET postulates the existence of a hierarchy of BVSR processes governing the scientific process, it has a clear normative role to play.

The aim of this section was to show that the general validity of the BVSR model can be defended. If this aim can be achieved, we will have at least a general argument to defend EET and its evolutionary agenda. Contrary to what he sought for many years, Campbell came to believe that this general aim cannot be achieved. Of course Campbell was partially right in arguing that the real issue regards the empirical question of whether the social BVSR mechanisms that EET postulates are aim-conducive in the actual. I fully agree with this but I also claim that the BVSR is epistemologically relevant. By endorsing such a pessimistic view Campbell gave up any interesting evolutionary reason to justify science, focusing instead solely on its social structure. But a mere sociology of scientific validity can easily give up its evolutionary commitments. In this section I have tried to defend Campbell's original agenda.

5.2.3 - EET as a sociology of scientific validity

In this subsection I will show that EET as a sociology of science delivers a social theory of the organisation of science, which provides a naturalistic account of the way science's social system is organised and how such organisation contributes to the achievement of science's goals. Science, from EET's perspective, is a functional system, partly adapted to the realisation of its institutional aims.

EET theorists aim to identify the BVSR mechanisms for the production of intellectual variants, their selection and preservation operative in science (e.g. Hull – cf. section 4.2 – and the later Campbell – cf. sections 3.3, 3.4 and 4.8 in particular). The relevant mechanisms are

reverse engineering.

social. EET aims to show that the postulated social BVSR processes contribute to improving our picture of reality.

In this section I shall try to articulate, by providing a few illustrative cases, what EET considers to be the sociological conditions for the validation of knowledge. In particular, such conditions pertain to the optimal organisation and maintenance of the vehicle carrying knowledge, where optimality is an ideal state. EET's claim is that the fulfilment of these sociological requirements is necessary to obtain and validate scientific knowledge.

EET is a naturalistic theory of science and as such it has an explanatory strategy that differs from that of other epistemologies. In two respects in particular this difference has to be stressed.

First, EET's explanations are causal. Traditional epistemologies usually provide rationalistic explications, which are the outcome of the rational reconstruction of the scientific events under analysis. From the EET perspective, however, the term 'rational' refers to ideal norms and not to causes of behaviour (Campbell 1997 p. 11), where ideal norms are explanatorily irrelevant from the naturalistic point if view. From EET's perspective, the causal factors of relevance are usually social. There is thus a sense in which EET's explanations can be assimilated to those proposed by other sociologists of knowledge (cf. section 5.4). EET's explanations focus largely on social elements and do not only refer to individualistic factors. More in particular, EET's explanations focus on the causal role played by the adaptive and epistemically conducive features of the social system of science. As we have already seen (cf. section 5.1) EET's analysis is populational rather than individualistic. EET's important causal claim is that at least sometimes the postulated BVSR social mechanisms are causally primary vis a vis individual level ones. This means that at least sometimes the individual selective decisions are directed by the social mechanisms (cf. section 5.3). In these respects EET's explanatory strategy is completely different from that of epistemologies that are not social and non-selective.

Secondly, EET studies the scientific process by reference to its social carrier or vehicle, whether this vehicle is a group of scientists or a more comprehensive scientific community. EET starts from the insight that, like any other social system and belief-preserving traditions, science must meet structural requirements in order to pursue its institutional aims. What is important to stress is that the requirements of vehicle or tribal maintenance act as co-selectors in the scientific process. Sometimes they act as assets, at other times as liabilities, interfering with the selective role that more relevant factors must play in order to have a successful science. In the latter sense many selective processes operative in science do not contribute to the validity of scientific knowledge. It is for this reason that EET must identify those functional (and adaptive) features of science that act as assets, since only these render science different from other institutions or self-perpetuating belief systems. Only in this way we can identify the selective features of the scientific process that more properly pertain to science and that justify a belief in its epistemic supremacy.[161]

We have already seen that EET theorists differ in the postulations they propose. For instance, in a rather Mertonian style, Campbell (cf. section 4.8) argues in favour of the causal role played by the anti-tribal norms of science, which, he continues, also explain why science is the most successful epistemic practice we have at our disposal. Instead, Hull puts special emphasis on the self-policing arrangements at the basis of the practice of science (e.g. its credit system, its reward and punishment system, its mechanisms for promoting mutual checking). More generally, EET distinguishes both between social and individual level BVSR processes, and between BVSR processes operating at the methodological level and those that work via social norms and institutional arrangements. EET puts special emphasis on the social mechanisms that work because of the causal influence of social norms and other institutional arrangements. I will give a few illustrations of these in the following pages.

The main explanatory aim at the core of EET's normative agenda is to differentiate conducive and non-conducive social BVSR processes.

[161] EET's solution to the demarcation problem will be treated in the next section.

The complexity of EET's descriptive and explanatory analysis presented in chapters 3 and 4 renders its normative agenda potentially very rich. If EET's analysis of the scientific process shows that some social BVSR processes have a beneficial role in the pursuit of the institutional aims of science (e.g. the achievement of a genuine community consensus), then the conducive causal role of these mechanisms must be protected. On the other hand, if the analysis shows that a social feature has a continuous deleterious role then recommendations ensue in order to curb its negative influence. EET's general normative aim is to suggest recommendations apt at either curbing or enhancing the influence of the social selective mechanisms that are respectively antagonistic and conducive to epistemic success and genuine consensus.

Before supplying some illustrations of the valuable contribution EET makes to normative epistemology, I have to highlight some of its limits.

The first pertains to the difficulty of identifying the conducive mechanisms of the social structure of science. The reason for this difficulty is that many adaptive features of the social system of science have evolved to play different functions in science. Sometimes they act as assets, sometimes as liabilities. The causal role played by a certain social feature is thus more or less dependent on the nature of the environmental conditions. For instance, even though a certain level of competition between individual scientists is necessary to have a successful science, it can also play a deleterious role (e.g. when papers are published well before being ready, thus carrying incomplete or even erroneous information). EET should therefore concentrate on those mechanisms that play a continuous and constant beneficial effect. I believe that all the examples I provide in the remainder of the section refer to continuously conducive social processes.

Secondly, EET provides recommendations aimed at either protecting the adapted features of the scientific process or at re-establishing and improving the effectiveness of the various selective social processes that render science a successful enterprise. However, given that from EET's perspective science is a functional system already partly adapted to the realisation of its institutional aims, its recommendations might be unneeded (apart from being difficult to enforce). This critical point highlights a difference in emphasis between EET theorists. In fact, while Campbell was an interventionist (e.g. calling for enforcing more respect for the anti-tribal social norms of science), Hull is a protectionist of what he considers an amazingly adapted social process. However, what is important to stress is that, despite their difficult enforcement, these recommendations remain normative. EET could be paradoxically an instrumentally useless epistemology, but for such purely philosophical purposes as mine this limit (yet to be established) remains irrelevant.

One of the most important social features of science is its demic structure, that is, its organization in research groups. Scientific research proceeds in demes. My claim is that the demic structure is an adaptive feature of the scientific process and that it is a necessary sociological condition for the validation of knowledge. As a consequence, the demic structure of science is one of those adapted features of science that EET recommends to protect.

More in particular, the advantages provided by a demic organisation of the scientific process pertain to the three levels of the scientific process characterised by the BVSR model.

At the variational level, working in a research group means being able to pool cognitive and conceptual resources. The group also provides an important closed environment in which to exchange not yet fully formulated ideas. The group also allows the possibility to engage in heated theoretical debates that would not be possible otherwise. All this means that the production of plausible hypotheses is enhanced by group activity and within group cooperation. As a matter of fact, the major vehicle of conceptual change in science is the research group.

At the selective level, the research group provides the principal mechanism for mutual monitoring. It encourages the severe testing of hypotheses proposed at the intra-group and inter-group level. At the intra-group level the selective process is generally more amicable and confidential. It mainly concerns the stages previous to the publication of results. At the inter-group level the process is generally competitive and public. The motivation to falsify a competitor's hypothesis is a basic feature of scientific practice that is fostered by the demic structure of science. More generally, the demic structure allows an adequate evaluation of the plausible hypotheses proposed both within and between groups.

At the preservation level, the group dynamics and the building of alliances increase the chance of ideas spreading. Isolated individuals would not be as effective in communicating and preserving memes.

My conlusion is that the demic structure is an essential social element of scientific practice because it promotes the achievement of the satisfactory levels of variation and selection that allow the BVSR processes operative in science to work conducively.

If this is the case then we can find evidence for my claim by considering cases in which the demic structure of science has had such alleged positive effects. I think that one such case is exemplified by the cold fusion polemics. Let me first briefly present the test case and then move back to my point.

After Pons and Fleischmann publicly presented their results on cold fusion, physicists all over the world attempted to replicate their experiments. In the United States the Department of Energy immediately set up a panel of prestigious scientists in order to test the claims made by them. Too quickly, the Energy Research Advisory Board report was rushed to publish. Its outcome was that no evidence of cold fusion was found in the further experiments undertaken and that Pons and Fleischmann's research was biased. It is quite interesting that the emergence of consensus in the community of physicists, namely that cold fusion is impossible, was reached despite future research (publicly but significantly "secretly" funded by the U.S. Navy)[162] partly vindicated Pons and Fleischmann's results. In fact, anomalous heat generation was recorded in many experiments, which even displayed evidence of nuclear reactions. Typical problems concerning the interpretation of experimental evidence ensued, because it still cannot be predicted adequately how much excess heat will be generated by using the same materials in the same experimental situation, and also because there is no consensus as to how

[162] The U.S. Navy was forced to fund such research by taking the money from "miscellaneous funds".

deuterium atoms overcome their natural repulsion, given that the results of the experiments on cold fusion carried out by U.S. Navy scientists do not fit with accepted physical theories. Factors such as insufficient evidence, difficulty in the replication of successful experiments, problems concerning the interpretation of available data, the need to revise accepted (but, I add, significantly "anomalous") theoretical knowledge, and consequent impossibility, more generally, of eliminating alternative explanations of the phenomena, should not provide "good reasons" to settle the dispute and declare cold fusion dead and buried. But this is exactly what happened. Consensus was not reached fairly because of the influence of social forces (e.g. influence of the petroleum-backed industry), even internal to the scientific community (e.g. envious physicists dreading the lack of public funds), which led to some bad science. In this case, the general conservativeness and competitive nature of the scientific process certainly had a dysfunctional effect in leading to the emergence of consensus. The consensus that cold fusion is nonsense resulted in "a breakdown in the process of unbiased, objective reporting of scientific information", according to Scott Chubb of the Naval Research Laboratory in Washington DC, guest editor of the journal "Accountability in Research".[163] In the cold fusion case the scientific process did not operate by approximating the ideal optimal strategy of investigation sketched in the previous subsection.

However, the fact remains that a relatively independent institution like the U.S. Navy decided to continue to do research despite the hostility of the community of physicists. One small research group constituted of well-qualified scientists funded by the U.S. Navy was almost secretly allowed to keep a plausible hypothesis alive and to pursue a proper evaluation of the not so incredible claims made by Pons and Fleischmann. A heresy was allowed to thrive, I claim, because of the way science is organized. If science was organised in a completely centralised manner, if science were not a competitive process, if science did not have the unintentionally designed evolutionary social features it has, then cold fusion would indeed be long dead and buried. It is clear that in this case the demic structure of science has had a

[163] "Accountability in research" October 2000, introduction.

positive effect on the progress of science, and this claim is true even if, as a matter of fact, the cold fusion hypothesis will be eventually eliminated for good evidential reasons.

What does EET recommend in similar cases of "normalisation" of scientific practice? My opinion is that EET should recommend defending the social arrangements that allow almost isolated research groups to pursue highly qualified research in crucial fields where inter-group warfare and competition is so high.

I also think that in this case one of the fundamental anti-tribal social norms of science was violated. I refer to the contravention of the anti-traditional norm of science. This norm has the function of injecting innovation and anti-conformism in a process that is fundamentally conservative like science. This norm is instrumental in challenging the tradition, which should be seen as a burden and a source of error rather than a source of revelational knowledge. In this case, those authoritative and powerful physicists who refused to accept limited but uncontrovertible evidence in favour of cold fusion were guilty of contravening such norm, to locate truth in some heroic past (the old, long-accepted and solidified physical theory), and to act in a tribe-like manner by defending the status quo of physical theory. In this sense the cold fusion case also shows that some curbing of group interests promotes knowledge validation.

EET also recommends to keep decentralised the management of the scientific process and to allow complete freedom of research, without political interference, to scientific institutions.

Another illustration of EET's normative approach can be presented by considering the population requirements relative to the scientific community. We have just seen an example of particular social structural requirement needed in order to render the scientific process conducive. But EET theorists also investigate the role that the "quantity" of people working in an area can have on the scientific process, rather than the ways in which these people are organised. In particular, EET theorists believe that a sufficiently big population of scientists working in a particular area is a necessary condition for achieving the "friendly competition" or "mutual exploitation" regime that allows the proper evaluation and elimination of hypotheses. Campbell (1997 p. 24-5) referred to this populational requirement as "critical mass". He argued that a critical mass is a basic sociological condition for the maintenance of "communities of truth-seekers". I share Campbell's insight. My claim is that critical mass is a necessary social condition for the achievement of epistemic adaptation because the BVSR processes operative in science can be conducive only given a sufficiently big number of scientists collaborating and competing in a field of enquiry. More particularly, a critical mass is necessary for the maximisation of relevant intellectual variation and for the proper evaluation and preservation of knowledge.

As far as the variational level is concerned, note that one of the few correct strict analogies between organic evolution and science as a selection process is that BVSR processes can produce adaptation only if selection can pick (by eliminating all the alternatives) the most interesting variants present in the environment. Thus, it seems almost obvious that, from the EET perspective, the larger the number of intellectual variants present in the intellectual pool, the more probable "stumbling" on an adaptive solution to a problem will be. As Campbell puts it:

> "The variations are, to be sure, bound to be restricted. But the wider the range of variations, the more likely a novel solution. The recommendation to speculate wildly thus belongs in the guide book to the strategy of discovery, if not in the logic."
>
> D. Campbell (1974b) p. 153

This principle points to the existence of an essential environmental condition for the emergence of knowledge that looks like a logical principle (cf. section 5.2.2). From a logical point of view, a multitude of variants will render the process of "stumbling" upon a good solution to a particular environmental problem more likely than a limited population – at least, as long as the amount of variation does not become so great as to overwhelm the selection mechanism. This essential environmental condition must be satisfied in the scientific case too.

A way to maximise the creation of intellectual variants concerns the nature of the scientific community. EET can study what are the optimal populational requirements necessary to preserve and optimise the potentially adaptive effect of intellectual variation. The populational analysis of the scientific process raises new issues. For instance: do we need a large population of scientists in order to maximise the possibility to produce innovative solutions to problems? How large must a population be in order to maximise the chance of "stumbling" on a problem solution? As I see it, the two basic alternative hypotheses concerning the populational structure of the scientific community are, first, that big populations are necessary to maximise the production of knowledge and, second, that they are not because a limited number of well trained and creative scientists is sufficient (in anthropology this view has been labelled the "theory of great minds"). These matters are to be assessed empirically and not a priori. Even though from a strictly logical point of view we could stretch the principle so as to argue that for an infinite population omniscience would be inevitable, we are interested in empirical and assessable claims. From the EET perspective we can say that there is evidence that the existence of a large population seems to be in some cases a necessary condition for the emergence of adaptation. There seems to be empirical evidence favouring the hypothesis that large populations are essential for the emergence and preservation of technological innovation (Diamond 1998).[164] But whether this evidence is directly relevant in the scientific case is a different matter.[165] The important point is, however, that the sociological conditions for the introduction and production of intellectual variation can be studied empirically with the aim to improve the epistemic conduciveness of the scientific process, and that recommendations can be derived from such empirical knowledge.

As far as the selective level is concerned, a critical mass is essential for adequate mutual monitoring, where the mutual check of each other's results is at the basis of the dynamic of inter-group competition in scientific practice. The critical mass is therefore an essential social element of the testing process and of the inevitably suboptimal evaluation of plausible

[164] Diamond (1998) argues that the best explanation of why a scientific culture emerged in the Middle East and Europe is not the theory of great minds or some kind of racist hypothesis, but rather a populational one. The same kind of explanation is used to explain the cultural poverty of, for instance, Tasmanian aborigenal culture.

[165] For instance, Hull (1988a p. 362 ff), who has studied this topic in detail, claims that in science a critical mass seems to be unnecessary because science is an elitist practice where most of the intellectual contribution is produced, at least officially (i.e. published papers) by few people. Note that the empirical evidence Hull reports would embarrass EET only if it were pursuing the analogical agenda.

alternative hypotheses. A critical mass is also necessary for preventing sectarianism by stimulating inter-group criticism and competition, and by contributing to the improvement of the communication between isolated research groups. In turn, better communication channels could also lead to the recombination of ideas sometimes necessary to overcome impending conceptual problems. The cascading causal effects of the critical mass requirement thus affect the typically interpenetrating variational and selective levels of the scientific process alike.

In particular, the survival or emergence of a scientific discipline depends on the ability of its advocates to recruit new members and to convince old enemies to switch allegiance. Until a critical mass is reached, the new research programme has no chance to survive long enough. A case to illustrate the causal role of the critical mass is the dispute about the existence of group selection. This is another case where community consensus was reached unfairly and where the scientific process worked sloppily, at least according to Sober and Wilson (1998 esp. chapter 2). Group selection was judged to be impossible too soon, and even after forty years of debate, with genuine evidence available that group selection exists in nature, the biological community is stuck to the idea that natural selection only works at the genetic and individual level.[166]

The reasons for the failed breakthrough of group selectionism are certainly many. One reason is the difficulty to "see" group selection at work, especially given that the predictions obtained by using individual-level and group-level models are many times undistinguishable. In particular, from the individual perspective group phenomena can be explained away without too much effort. Also, the two models explain equally well a vast range of biological phenomena and the decision to prefer one model to the other is in such cases purely pragmatical.

A further reason concerns the selective role played by metaphysical beliefs on the scientific process, and how this role is sometimes even greater than that of empirical evidence. In this sense, one reason why group selection was declared dead forty years ago, and certainly the

[166] A usually "liberal" scientific journal like the "New Scientist" (15.3.2003), in its section (159) "Inside Science" about the evolution of cooperation, states that "One suggestion – now discredited – was that selection acts not on individuals but on groups of organisms."

reason why it is still difficult to resuscitate it, is determined by the appeal reductionist (i.e. individualist) thinking plays in any community of scientists that works to render their discipline respectable and "scientific".[167]

But a further reason is sociological and pertains to the inability of group selectionists to build a critical mass. It is quite symptomatic to know that trainees in biology were, and still are, reminded that group selection is an heresy, a subject not even worth a tentative look, exploration, study (Wilson 1989), where the trainees' acritical compulsion to follow the "authorities" in the field could be interpreted, following Hull, as the attempt to maximise their conceptual fitness. Unfortunately, no institution like the U.S. Navy could come to rescue in this case, given that the technological application of group selectionist ideas seems less promising than the possibility to extract potentially unlimited amounts of energy from sea water. Of course it is difficult to test the counterfactual claim that if trainees in the biological sciences were not disillusioned and brainwashed to swear allegiance to the individual selection paradigm the outcome would have been different. However, EET's contention can be vindicated by applying its recommendations in similar cases.

My recommendation would not be to defend the critical mass requirement uncritically. I do not think that Biblical scholarship should be allowed to thrive in this sense. History will take its course and recruiting new members will be left in the hands of the practitioners. Contingency plays such an important role in science (analogously to what happens in biological evolution). However, it is difficult to resist endorsing a variant of Plank's principle at this juncture. Planck famously said that:

"...a new scientific truth does not triumph by convincing its opponents and making them see the light, but rather because its opponents eventually die and a new generation grows up that is familiar with it."[168]

[167] Individualism (both as a methodological and ontological thesis) has become the mark of the scientific; it works in science as an unquestioned substructure of thought, as an unsubstantiated assumption concerning the proper way to conduct scientific research. From an individualistic perspective group properties and group selective processes do not exist. In section 5.1 I showed that despite the fact that individualism still is the mark of respectability for scientific and epistemological research EET does not endorse individualism.

[168] Taken from Hull 2001d p. 226.

The variant of the principle I am interested in is that good hypotheses can be lost along the way because of populational reasons. If this is the case (and empirical studies are needed to ascertain if it is) then critical mass is a necessary condition for improving the conduciveness of science. As a consequence it is a social feature of science that should be protected. At the variational level, a critical mass ultimately results in an enhanced possibility to formulate fitter hypotheses by increasing the chances of within-group and between-group memetic recombination, by encouraging between-group communication and interbreeding, and more generally by injecting innovation in a conservative process where the actors refuse to adequately evaluate putative plausible hypotheses. On the selective side, what matters is the overall population in the area of interest rather than the number of scientists working on the heretical project. The former should not be allowed to become too small. The populational requirement is a necessary condition for improving the quality of the testing process in scientific practice, a process that crucially depends on the mutual checking between competing groups. But in order to check other groups' results you need to see the competing research group as a group. And without a critical mass this cannot be done. Scientists working in complete isolation can certainly produce high quality work, but working in a group increases the chances of being visible (e.g. by having access to more publications). It would be interesting to study the recent history of science and check how many plausible hypotheses were publicly proposed by completely isolated and detached scientists, and verify how many of them have survived. My prediction is that the number is very small.

How to realise the recommendation in practice is more difficult. To say that the recruiting process should be unbiased and fair is a true enough platitude. But how can this recommendation be enforced? To encourage a more critical attitude towards the accepted theoretical status quo by preaching, as Campbell argues, the anti-tribal norms of science more effectively is also problematic: scientists are not going to accept this kind of advice from EET theorists. This means that recommendation enforcement might be useless. However, this should not be an argument against EET, at least for the reason that even what more traditional epistemologists preach is completely neglected by scientists. Here we have a contrast between the two attitudes I prefigured at the start of the section: interventionism and laissez faire. My point is that Campbell's interventionism might be as conducive epistemically as Bush's military interventionsm in Iraq is conducive to world peace. EET can devise and recommend policy solutions of various kind in order to increase the levels of variation in science, in order to maximise the number of plausible ideas generated, in order to render more severe the process of mutual monitoring, in order to allow the more objective and less partisan evaluation of the plausible hypotheses, all this with the ultimate normative aim to improve the epistemic conduciveness of the scientific process. The possibility to come up with recommendations is more than enough to show that EET has a normative value. But recommendation enforcement is a completely different matter. Hull has repeatedly pointed out that the scientific process is already largely adapted. This means that recommendations are unnecessary most of the time because the social system will reach a new level of adaptedness without external intervention (which is likely to be detrimental most of the time). In this sense Hull's laissez faire attitude is in line with his slim normative approach.

Hull's position can be understood by referring to those structural arrangements of the practice of science that allow, for instance, the introduction of intellectual variation in a process like science. Some of these arrangements are far from artificial or superimposed, but are rather "natural", that is, part and parcel of the invisible hand at the heart of the scientific game posited by Hull, or, less mysteriously, part and parcel of the evolutionarily designed social structure of science. Consider as an example the idiosyncratic ways scientists use to interpret available evidence. This phenomenon has been described but, arguably, never properly explained. I think EET can help in this respect.

It is generally assumed that scientists should interpret available evidence in the same way. This is because, according to traditional philosophy of science, there must be one correct way to interpret the evidence. Naturalised epistemologies like EET are, for reasons that should be obvious by now, not keen to take such a fundamentalist stance. If the description of scientific practice inevitably shows that this is how scientists behave, and if science continues to be successful, then there must be a reason why such behaviour is not detrimental to the achievement of the institutional aims of science. Naturalists are thus engaged in understanding the function of the behaviour rather than in showing that such behaviour, by being incompatible with idealised models of rational action, should be called irrational.

From the EET perspective, the deviant and idiosyncratic behaviour of scientists in dealing with the always finite available evidence has a clear function. Kuhn was the first to individuate this function when he stressed that a selective decision based on a fixed methodology (e.g. an algorithm) would produce a uniform choice that would affect the epistemic conduciveness of the scientific process. The fact that scientific method is not a static system of rules as traditionally contended but rather a value system, Kuhn argued, has thus a positive causal effect. This is because scientists are not forced to behave in rule-following manners but are rather allowed to apply and adopt the methodological canons (accuracy, simplicity, scope etc.) differently.[169] Hence, Kuhn argued, behavioural variability in scientific selection has a positive functional role because:

"With standards set higher, no one satisfying the criterion of rationality would be inclined to try out the new theory, to articulate it in ways which showed its fruitfulness or displayed its accuracy and scope. I doubt that science would survive that change. What from one viewpoint may seem the looseness and imperfection of choice criteria conceived as rules may, when the same criteria are seen as values, appear an indispensable means of spreading the risk which the introduction or support or novelty always entails."
T. Kuhn (1977) p. 332

I think that EET explains in novel terms this well-known fact about scientific selection. The idiosyncratic behaviour of scientists in matters of selection is just a "natural" and bottom-up mechanism designed to introduce significant intellectual variants. The difficulty to see this point for traditional epistemologies depends on their individualistic perspective, from which

[169] First, evaluative criteria can be applied to real problem situations differently by different scientists. Secondly, the criteria, when employed together, might conflict with each other, that is, following one of the criteria with extreme loyalty will necessarily interfere with compliance with the others.

scientists' behaviour, when not rule-abiding, looks blatantly irrational and unjustified (cf. section 4.4). However, from the EET populational perspective it can be interpreted as an instance of the non-enforced processes aimed at securing the adaptedness of the scientific process Hull refers to. More specifically, the mechanism supplies a means to achieve the aim of maximising the number of plausible hypotheses formulated by approximating the optimum prescribed by the optimality BVSR process analysed in section 5.2.2.

The three illustrations I presented in this section show that EET has a clear normative agenda. The three examples I provided are significant because they vindicate EET's thesis that the social structure of science contributes to the validation of scientific knowledge. By identifying some of the necessary social conditions for the emergence of scientific knowledge, EET theorists can try to devise policy solutions aimed at protecting the relevant social features of the scientific process. In this sense EET fulfils the meliorative goal central to epistemic naturalism. I thus believe to have shown that EET yields a sociology of scientific validity as contended by Campbell and Hull.

I would like to stress an important point concerning the nature of the recommendations produced by EET. EET as a sociology of science shows that neglected social factors about the organisational structure of science are epistemologically relevant. The social BVSR processes operative in science are social norms, various kinds of institutional arrangements, and furthermore social maintenance requirements.

A crucial point is that in the social structure of science are already included the requirements about testing. This inclusion makes EET's and traditional epistemologies' recommendations sometimes similar, specifically when the methodological aspect of the scientific process is causally most relevant. However, at the methodological level selection is an individualistic matter: it is most importantly individual scientists who apply, for instance, modus tollens. This means that EET has a positive contribution to give to epistemological analysis mainly as far as the social level is concerned. At the individualistic level what EET recommends is largely consonant to what other traditional epistemologies suggest (cf. section 5.3). This is because EET's populational account implicitly comprehends the methodological element (e.g. the selective role of the testing process). In this way EET's approach can be seen as not fully articulated.

If this is the case, where does the originality of EET's approach lie? The originality of the approach concerns its methodology and EET's contention that, at least sometimes, the social BVSR processes are causally primary. I believe that EET provides a genuinely original perspective in both senses. But the originality should not be stretched too far. In section 4.6 I criticised Hull's EET because his epistemological analysis regards merely the population level, the only level, he claims, at which fruitful generalisations concerning science as a process can be found. I repeat here the same point: EET's populational explanations of the success of science only identify part of the causal factors influencing science's success. This is because we can give explanations of the success of science from at least two levels of analysis. The first level concerns individual scientists and the application of methodological canons, the second concerns the population level and the social structure of science.

To conclude, the point I want to stress is that to emphasise the social perspective should not amount to neglect the methodological aspect of the scientific process. EET should not commit Hull's mistake. This is because of one main reason that will be considered more fully in the next section: the demarcation between science and pseudo-science cannot be achieved if the methodological aspect of the scientific process is not taken into account. I agree that, as Campbell and Hull contend, social norms and social features differentiate science from other epistemic practices. But this is only part of the answer.

I must also highlight one major limit of EET's normative agenda as proposed by Campbell and Hull that I wish to overcome with my characterisation.

The problem was already mentioned in the previous sub-section and has to do with the limited normative outlook afforded by a mere sociological perspective. Campbell and Hull believe that selection theory does not provide any justification of knowledge. It is quite paradoxical that these theorists, for years at the forefront in the battle to develop and defend the EET programme of evolutionary epistemology, finally recanted by admitting that the validity of scientific knowledge stems only from its social structure and not from the fact that it is a selection process. I think that this admission turns Campbell's and Hull's sociology of scientific validity into a kind of Mertonian sociology with a very indistinct evolutionary flavour. There are social norms (universalism, disinterest, communism and organised scepticism), and in addition there are institutional arrangements, and furthermore social maintenance requirements. But how do the evolutionary and sociological perspective combine? As a matter of fact they do not because the former is completely abandoned. My aim is to retrieve the evolutionary element and I think that this can be done for the reasons already highlighted in the previous section.

My claim is that both the explanations and recommendations provided by EET make sense in the context of the general theory of reliability and epistemic conduciveness outlined in section 5.2.2. There I argued that the BVSR model provides epistemological warrant because it characterises a naturalistic approximation to the ideal and optimal method of investigation. In this section I tried to show that the BVSR social processes posited by EET are reliable because they are conducive to approximating the optimal investigative strategy of eliminating the maximum number of plausible alternative hypotheses. I showed, by giving three illustrations, that EET tries to identify the institutional arrangements, social norms and sociological conditions that permit us to approximate the optimal levels of intellectual production, selective evaluation and preservation of ideas and hypotheses. The recommendations I proposed were also to be understood in the same terms, that is, as recipes and suggestions aimed at protecting the social features that allow the approximation of the optimum.

5.3 Evolutionary constructivism

In section 1.2 I pointed out that some form of relativism is a necessary outcome of the naturalistic turn. In fact, some form of relativism follows from the naturalist's choice of reframing the normative project as an instrumental question about the improvement of our cognitive performance: given that naturalistic norms are hypothetical and dependent on contingent hypotheses about humans, then, as a matter of fact, the scientifically extrapolated epistemic principles necessary to improve cognitive performance in the actual world might turn out to be far from universal.

The lack of universalism of the naturalistic approach can be assimilated to the thesis I called, following Stich, normative cognitive pluralism. This thesis states that there is no unique set of standards of epistemic evaluation that epistemic agents ought to use, but that such standards are relative to vagaries of contingent facts (e.g. about an individual, about a group of epistemic agents, about specific epistemic circumstances). In the same section I also pointed out that the alternative to normative cognitive pluralism is normative cognitive monism, that is, the much stronger universalistic thesis according to which there exists a single set of standards that epistemic agents ought to use. It is important to stress two points at this juncture. First, while normative cognitive monism is a monolithic thesis, the thesis of normative cognitive pluralism comes in degrees. Normative cognitive pluralism comes in many forms, some more relativistic than others. The spectrum of possibilities varies from radical relativism to much weaker forms. I suggest that this is the reason why the term "relativism" is so difficult to define epistemically (i.e. it refers to many theses of various strength). Secondly, not all forms of normative cognitive pluralism are dangerous (e.g inevitably yielding vicious kinds of cultural relativism). This is exactly what I am going to contend in this section, where I will show that EET is committed to a limited and non-dangerous form of normative cognitive pluralism.

The serious, dangerous and threatening kinds of epistemic relativism concern the thesis that the standards of epistemic evaluation are relative to facts about groups of epistemic agents, that they are, in brief, community-dependent. The putative risks connected with such kind of relativism are various. First of all, relativistic standards might be so sensitive to facts about the community of interest as to become non-applicable across cultures and contexts, so that it will be impossible to assess claims of epistemic practices that do not share standards in common with our practice. In this way we cannot discriminate between science and pseudo-science, since pseudo-scientists will always argue for the existence of a different set of standards regulating their practice. Secondly, the relativistic standards of epistemic evaluation might turn out to be insufficient to "rationally" solve many cases of inter-cultural methodological and theoretical disagreements. If this is the case then scientific change can be properly seen as an arbitrary and unregulated process, always (or at least more often than not) underdetermined by good reasons. At this point the aim I have in this section is more delineated. What I need to assess is the following:

1) Does EET extrapolate useful naturalistic norms general enough to be applied to different contexts?
2) Does EET provide means to solve the demarcation problem?
3) Does EET offer any resources to help solve inter-theoretical debates?
4) Does EET offer any good reasons to select and prefer certain hypotheses over others?

EET's first contribution to rebut relativism concerns the articulation of the notion of epistemic norm or evaluative standard. Epistemic standards can be seen, from the evolutionary perspective, as evolutionary products (i.e. product of a selection process). In science such processes are not biological, but cultural and social (cf. section 3.2). Epistemic standards are thus generally cultural and social products and can be studied by focusing on the details of their cultural emergence and especially on the role they play in the scientific process. The norms of science constitute a group of cultural artefacts culturally transmitted and "evolving" by cultural selection. From an historical point of view, the methodological and institutional norms of science are evolutionary products and adaptations.

The advantage of seeing norms as evolutionary products arises from their being treated as historical objects, where such objects do not belong to a particular perspective. Rather, the epistemic standards studied by EET have an autonomous life and a community-transcendent nature. In this way the objectivity of science is not dependent on the judgement, for instance, of a particular community of trained scientists, but it is constrained by the causal role that some multi-generational standards play in scientific practice. As Campbell put it (1981 p. 514), the objectivity of science depends on the disputations between "multigenerational communities of truth-seekers".
To summarise, from a metaphysical point of view, EET's thesis is constructivist. Epistemic standards undergo a process of selection in which many generations of the scientific community actively participate (in the sense that subsequent communities can re-assess and revise the standards). Epistemic standards can be both biological and cultural adaptations.

Leaving the metaphysical issue on the side for the moment, the details of EET's rebuttal of dangerous kinds of relativism should be judged by assessing the epistemic value of the extrapolated standards. In this sense we can distinguish between two kinds of norms captured by EET: the methodological and the social ones. An important point is that, prima facie, what EET has to say about the methodological aspect of the scientific process is largely consonant with what many traditional epistemologists would argue, namely that the crux of scientific method is to eliminate rival hypotheses through empirical testing (cf. section 5.2.2). In this sense, Campbell claimed that spelling the details of how a selectionist model 'justifies' scientific knowledge would "lead to rather orthodox conclusions" (Campbell 1997 p.7). This is a further way to argue for the complementarity of traditional and evolutionary epistemologies that should not, however, be stretched too far (cf. section 5.4). In fact:

"From my perspective the ideology and norms of science are not clearly distinguished from 'scientific method'. Scientific method is also to be seen as a product of cultural-evolutionary process on the part of a bounded belief-transmitting subsociety of many generations.... While historically both methods and ideology have fed on concrete successes, it is convenient to regard the ideology and practice of cooperative truth-seeking as coming first and method as a rationalized summary of successful usage in the community."

D.T. Campbell (1981) p. 514-5

EET is thus committed to view the social norms as causally primary and the methodological ones as secondary, contrary to what preached by traditional epistemologies.

As we have seen in the previous pages, both Hull and Campbell argue that what is needed in order to have a successful science is an adaptive social structure rather than a community of methodology-abiding scientists. The important point stressed by EET theorists is that the social structure is primary in explaining scientific success. On the one hand, Campbell emphasised the causal role of the social norms of science, which, as evolutionary products, have been unintentionally designed, at least, in the course of the 300 years of official history of science by various selective forces in order to channel consensus and maximise the selective role of nature in the editing of our picture of the world (cf. section 4.8). On the other, Hull argues that if the institutional arrangements regulating the scientific game are functionally organised then individual scientists' behaviour becomes somehow irrelevant because it will be the invisible-hand (cf. section 4.6) that will render the process conducive. Despite differences in emphasis, their points are the same.

The point I am stressing further explains why the social structure of science plays a normative role (cf. section 5.2), and why EET's methodology must be populational (cf. section 5.1). EET's big idea is that individual scientists' behaviour is causally secondary, somehow epiphenomenal, because evolutionary designed selective mechanisms (i.e. some examples where provided in section 5.2.3) are operative that force individual scientists to behave in particular ways despite their epistemic motivations. This also explains why, in practice, the selective behaviour of scientists can be "safely" idiosyncratic without affecting the conduciveness of the process. This is because, again and more technically, the vicarious social selective mechanisms realising the self-policing system of science take the place of rule-abiding methodological winnowing at the individual level.

But, the emphasis on the social rather than methodological nature of scientific practice does not amount to a rejection of the latter. To return to Campbell's illuminating passage, methodological norms can be easily assimilated in the EET framework.

General selection theory can provide an explanation of the workings of the evolutionary process of emergence of the norms of science, for instance by relying on the hypothesis of cultural group selection. This historical thesis potentially accounts for the emergence of various kinds of standards (e.g. Popperian regulative ideas – cf. Popper 1979 chapter 3, Campbell's anti-tribal norms of science, Kuhn's ideological commitments, Hull's institutional arrangements regulating the credit-checking system of science). But how does a hypothesis about the genesis and emergence of the norms of science validate them? In section 5.2.1 I have already explained that EET does not commit the naturalistic fallacy. It does not even commit the genetic fallacy. EET is committed to view norms as natural objects, which can only be tentatively validated via empirical means. Furthermore and crucially, from Campbell's passage above we also evince that EET sees the norms as inductive generalisations of proved reliability, in a rather Millian fashion.

EET provides an original explanation of why science is a superior epistemic practice and of why it is sometimes successful in describing nature. Even if it is true that science has a lot in common with other epistemic practices and belief-preserving traditions (e.g. what Campbell calls its "tribalism"), it remains a vastly superior epistemic practice because its social structure is adapted and its institutional arrangements contribute, in normal circumstances, to the validation of scientific knowledge. In section 5.2.3 I have argued that there exist specific sociological features that are peculiar to science, that differentiate the practice of science from other tribe-like traditions, that render science superior to all the other epistemic practices, and that allow scientific knowledge to be validated independently of what a certain community of trained scientists believes.

However, EET's sociological solution to the demarcation problem could be criticised for the following reason. If the epistemic standards governing the scientific process are evolutionary products, and if the selective processes that led to the emergence of such norms are those that led to the rise of a particular historically emerged epistemic practice, namely contemporary science, then some kind of cultural relativism is inescapable.

Creationists could contend that the problem of EET's approach is that it is enslaved to a very specific sociological perspective: because science arose in the 17th century by challenging, for instance with Galileo, the authority of the Catholic Church, it exhibits certain sociological and methodological features like, for example, its anti-traditional and anti-authoritarian social norms; however, the creationist continues, the fact that creation science defends the authority of the Bible and does not endorse the social norms identified by EET is simply due to the contingencies of its development; but it would be wrong to conclude, creationists continue, that creation science is pseudo-scientific merely on the grounds that it lacks the sociological features of contemporary science.

Something appears to be too relativistic in the sociological solution of the demarcation problem proposed by EET because the social and methodological norms to which EET refers identify properties dependent on the peculiarity of the social vehicle that is contemporary science. EET's selective criteria thus are far from being objective, but are rather relative to the environmental situation in which science has evolved and in which we now play the scientific game. But if this is the case, then the objectivity of science really becomes just a matter of persuading and converting members belonging to a particular epistemic traditions (e.g. creationism) to another (e.g. Darwinism).

I believe this argument is not threatening for EET. First of all, EET relies on certain methodological criteria of selection (e.g. modus ponens or the method of comparative testing), whose reliability is not under discussion. In general, the norms that EET identifies are of presumptive reliability, being inductive generalisations summarising, as Campbell put it, past successful usage. Secondly, and more importantly, the original claim made by EET theorists is that the social norms and institutional arrangements of science are epistemically conducive. For this reason, any epistemic practice that is so organised as to reject, like creation science, the value of the anti-tribal norms at the core of science might be a well-structured belief-transmitting practice, but it will not be an epistemically conducive one. And this is again because the anti-tribal norms of science allow empirical evidence to play a significant selective role in weeding out bad and implausible hypotheses, while any social arrangement aimed at maximising the authority of the Holy Book does not so.

Finally, to the questions "Does EET offer any resources to help solve inter-theoretical debates?" and "Does EET offer any good reasons to select and prefer certain hypotheses over others?" the answer is that such reasons exists and that they are evidential and methodological, where these evidential and methodological standards are somehow universal (in the sense that their application across contexts is generally reliable). But the additional point EET makes is that, given the primary causal role of the social structure of science, such universal epistemic standards, as the analysis of scientific practice shows, are allowed to be applied by different scientists idiosyncratically both because selective variation has a functional role in science that seems to allow rather than hinder scientific progress (cf. section 5.2.3), and also because individual scientists' behaviour is, as illustrated in this section, "directed" by the evolutionary designed selective mechanisms operative at the social level. This explanation vindicates both EET's contention that the selective role of the social structure of science is necessary for the validation and accumulation of scientific knowledge, and the other contention that some good reasons for hypotheses selection can be said to be social rather than individualistic. In any case, to answer the questions above mentioned, EET is committed to the existence of various kinds of good reasons to justify hypotheses preference.

To conclude this section, I shall now characterise the way in which EET is committed to normative cognitive pluralism. EET advocates a thesis that can be referred to as "perspectival relativism" (Campbell DE pp. 449-450).[170] Perspectival relativism states that knowledge is always partial, incomplete and relative to the limited phenomenal perspective of the knower.

[170] In order to characterise this position recall Campbell's definition of knowledge in terms of the fit between the representation and the represented (cf. section 3.3). Campbell used the metaphor of the map to explicate this definition: the map somehow fits the territory and for this reason it provides some kind of knowledge. The metaphor captures two features of knowledge. On the one hand maps are always partial and incomplete representations. On the other they describe some aspect of the territory of relevance to the map maker. Analogously, even though knowledge is inevitably representationally incomplete, it describes some aspect of reality of relevance to the knower.

There is only a view from somewhere, from a particular perspective, where this view is determined by the characteristics of the vehicle carrying knowledge. In the case of scientific knowledge we can consider at least two ways in which the vehicle can be characterised.

If as vehicle we consider the individual scientist then perspectival relativism states that knowledge is partially determined by the limited cognitive characteristics of the individual agent. The issue concerns the nature of the human cognitive strategies used in science. One important current debate amongst naturalists concerns the nature of the psychological norms that naturalised epistemologies are able to justify. In particular, the camp is divided between those who sustain the thesis of the species-typicality of such norms and those who defend the legitimacy of a stronger universalistic approach. I believe that the only thing that can be fruitfully said concerning the individualistic properties of the vehicle carrying scientific knowledge is that, science being a peculiar human cultural phenomenon, it is at this stage impossible to verify whether the ways in which we ought to edit our beliefs are more than species-typical.

If by vehicle we consider a social entity then perspectival relativism states that knowledge is partially determined by the social structure of science. Contemporary science certainly has features that it would not have had if the historical process that led to its emergence had been affected by different contingencies. In this sense, some of its methodological (e.g. the preference for theories only positing observables processes and entities) and sociological features (e.g. its competitive nature) might be just legacies of a particular evolutionary pattern.

Both the thesis of the species-typicality and the thesis of the historical contingency of the norms of science are concessions to the relativist as they affect the belief in the universality of the norms EET identifies. It is certainly true that EET norms are relative to facts about how science as a practice evolved in the course of modern history, to contingent facts about how science is structured, and to facts about us as a species capable of cultural evolution. But in the previous pages I have shown that all these contingent factors do not affect the possibility to extrapolate norms general enough to be applied across contexts. And this means that perspectival relativism is not a threatening form of relativism.

5.4 Hypothetical realism

In the previous pages I pointed out that what EET has to say about science is compatible with what many traditional epistemologists would argue. In section 5.2.2 I showed that EET shares with many epistemologies an eliminativist commitment, while in section 5.3 I showed how EET views some methodological and social standards of evaluation as applicable across contexts. This is further evidence for the thesis of the complementarity and peaceful co-existence of traditional and evolutionary epistemologies I have been arguing for (cf. chapter 1). On the other hand, in the rest of this chapter I have also tried to show that EET's explanations and recommendations are different from those produced by many alternative approaches. In this final section I will move further in this direction by showing how EET's position differs from social constructivism. Then I shall consider what its realist commitments are.

In section 5.3 I argued that EET relies on a constructivist thesis to save scientific objectivity: EET sees evaluative standards as evolutionary products that "transcend" our particular social perspective. But the term "constructivism" was used in a rather different way from how it is normally used in the epistemological literature, where the term refers to epistemologies generally advocating the thesis that knowledge is determined mainly or solely by social factors rather than empirical evidence. Social contructivists treat science as a manufacturing institution and challenge the ontological thesis at the heart of the doctrine of scientific realism, namely that there exists a mind-independent reality to which knowledge should refer to. Of course EET is committed as well to the view that social factors determine the content of knowledge, however indirectly and partially. The difference between the positions is that EET aims to identify and protect those mechanisms internal to the scientific process that contribute to the maximisation of the selective role of evidence.[171] This also means that EET is not committed to the anti-realist stance typical of social constructivism. Nonetheless, EET endorses some constructivist theses. In order to illustrate the similarities between the two epistemologies I will refer to Latour (1987)

I believe that the anthropological science studies undertaken by Latour have had an important critical function. In two respects Latour's contribution has to be stressed.

First, Latour shows that the acceptance of hypotheses is not a matter solely related to evidence but that it involves social negotiations. I believe that Latour's concept of black box is useful to understand this process. Scientists have to agree and reach consensus on how to understand the data, on how to interpret the available evidence, on how to model nature, and these processes involve a lot of theoretical discussion and competitive argumentation within and between research groups (Latour particularly stresses the laboratory level, but more comprehensive social levels have to be taken into account). When consensus is reached (typically, according to Latour, in a rather political fashion) hypotheses become crystallised and are never again re-assessed or challenged. I believe this is a real phenomenon in science that points to the sub-optimality of the testing process. This phenomenon has also been stressed by Hull. To identify the black boxes becomes, from the EET perspective, one of the aims of epistemology and science itself.

Secondly, Latour describes the social dimension of science and stresses the epistemic role played by the credit system (e.g. the practice of citing leading scientists to lend authority to one's otherwise empirically inadequate results) and by the funding system (e.g. the vicious networking process going on at the cutting-edge of research to get funding). All this is very familiar to EET theorists. One reason for this similarity in views might be partly due to the fact that Latour and EET theorists share the same methodology (i.e. methodological populationism).

I shall now briefly highlight the differences between constructivists like Latour and EET theorists. The first basic difference between the approaches is that for Latour (cf. 1990) science is not "a" process but a series of practices. Latour contends that an analysis of actual scientific practice shows that, for instance, the manufacture of knowledge in biology is

[171] This peculiar position can be called "internalised externalism". Internalised externalism resembles in many ways

completely different from the manufacture of physical knowledge. EET theorists disagree. A second basic difference already pointed out regards ontological matters. Despite typically endorsing a “methodological” constructivism (i.e. the constructive aspect of the scientific process is of epistemological relevance) Latour sometimes seems to commit his view to an unmotivated ontological constructivism (i.e. the world is entirely constructed by scientists). Latour’s claim that black boxes cannot be given an existence beyond their instrumental setting and laboratory environment seems to me an unjustified anti-realist commitment.

In brief, Latour emphasises the constructive aspect of the knowledge process, while EET theorists are more "sensible" (I believe) to argue for the co-selective role of evidence. As a matter of fact, the general coherence of the constructivist approach can be doubted: “Relativist students of science present numerous case studies to show how unimportant anything that might be termed evidence is in changing scientists’ minds, but something is desperately wrong with presenting evidence to show how irrelevant evidence actually is.” (Hull 2001d p. 239).

From these brief remarks we can conclude that even though EET treats knowledge as partially constructed by a social vehicle, evolutionary constructivism is a completely different thesis to social constructivism. The basic difference is that from EET’s perspective the methodological and social norms of science are adaptations.

I have so far stated that EET is not committed to an ontological constructivist thesis. However, I have not yet considered the nature of EET’s realist commitments. The issue of realism was introduced in section 1.1 where I distinguished between three related realist theses. The ontological thesis concerning the independent existence of the objects of scientific knowledge from the knowing subject, the epistemological thesis according to which scientific knowledge is somehow representationally true, and the axiological thesis that truth is the aim of enquiry. In the following pages I will suggest that while EET’s commitment to the first thesis is less problematic, its commitment to epistemological and axiological realism is more difficult to justify.

Merton’s classic sociology of science.

EET roots are apparently well planted in the realist tradition. EET is committed to a view that can be labelled hypothetical realism. In Campbell's (1977 p. 445) formulation hypothetical realism is the view that there exists an objective world independent of any knowing subject that scientific theories only approximately fit, and that truth is the aim of science. Even though we can never directly compare belief with reality but only belief with belief, Campbell retains the realist ontology. The existence of an external reality that knowledge aims to "fit" is thus an assumption.

How does EET justify this assumption? The only kind of argument I have found in the literature in evolutionary epistemology aims to show that ontological realism is a functionally fruitful posit. For instance, Campbell (1974a p. 450) argues that psychological evidence shows that many mammals have evolved to an epistemological state where they adopt a realistic hypothesis concerning the nature of space, that is, the various separate spaces classified according to utilitarian motives (e.g. thirst space, hunger space, mate-finding space) are assumed to be one all-purpose space. This is evidence, Campbell argues, in favour of the veracity of the dualism organism-environment at the core of the biological sciences. Unfortunately, Campbell's argument is not convincing. Even though it were shown that we are "programmed" to think in specific realist ways, for the radical sceptic or the social constructivist this would still not constitute evidence of the existence of some ontological posit.

How does EET answer the constructivist ontological challenge? My opinion is that a reply is not required. This is because the constructivist case is not on a better epistemic grounding. In fact, even if knowledge is partially constructed, I do not see any good argument to extend the constructivist argument to the extreme by arguing in favour of idealism. With this I do not mean to deny that the constructivist challenge is "partly" significant. Latour might be right in pointing out that black boxes are normally used in science, but to argue that external reality is a black box amounts to stretching a good point too far. The role of the epistemology of science is to get rid of false reifications. But to get rid of external reality would amount to abandon science, epistemology and knowledge altogether.

The positive case in favour of the endorsement of axiological realism is even dimmer. One reason is that evolutionary considerations are generally seen as directing us towards embracing a kind if instrumentalism. In fact, the aim of natural selection, if any, is fitness maximisation rather than truth, where fitness is measured in instrumental terms (i.e. survival and reproduction).

A reasonable argument in favour of axiological instrumentalism along these lines is proposed by Rosenberg (1996). We know (from section 1.2) that naturalistic epistemologies endorse scientism. We also know that the thesis of scientism needs a naturalistic justification. Rosenberg believes that this sort of justification can only come from the "fact" that contemporary science is more reliable than past science, where reliability is measured in terms of science's power of prediction and control.[172] So, if science delivers the goods in instrumental terms, this means that prediction and control are to be endorsed as the aims of science. In fact "Darwinian theory tells us that no one could long survive who does not embrace prediction as the aim of enquiry" (Rosenberg 1996 p. 27).

In EET's case this argument is not sufficient to reject truth as the aim of enquiry for two reasons. First, even though natural selection cares about fitness and predictive adequacy rather than truth, EET does not model scientific selection as an entirely analogous process to natural selection. The relevant issue is only whether the social selective processes EET posits are truth-conducive, not whether BVSR processes in general are. Secondly, instrumental reliability could be an observable measure of truth.

However, the fact remains that the rejection of the analogical reading must be coupled with substantive arguments aimed at showing that science can and should aim at true instead of

[172] Rosenberg's aim is also to show that the reliability of science provides a non-circular justification of scientific practice, solving the problem of circularity we treated in section 1.2. As a naturalist I agree with Rosenberg. Unfortunately non-naturalists would not, since for them the demonstration of the reliability of science is based on standards of evaluation that precede such demonstration. For non-naturalists, that is, the claim that science is instrumentally reliable is not a description of a fact, but is normative.

instrumentally reliable knowledge. Distancing the evolutionary approach from its instrumentalist underpinnings might be possible, but I don't know of any convincing arguments in favour of realism produced by EET theorists. The only argument proposed is that thinking like a realist is more productive than thinking like an instrumentalist and that a realist ideology retains a functional role in science. In brief, the commitment to axiological realism is enforced by the social system. In this sense, Campbell (1974a p.450) claims that the history of science shows that those scientists who interpreted their theories realistically "have repeatedly emerged in the main stream for future developments", while Hull (1992a) challenges anti-realists to test their philosophical claims against realist ones.

Hypothetical realism is also committed to scientific realism. But even this latter thesis is hard to justify from the EET perspective.

EET theorists are committed to the view that science progresses. In order to justify such a belief EET theorists have proposed different arguments. Hull believes that the social organisation of science is largely adapted to the achievement of true knowledge, even though some environmental conditions need to be satisfied for progress to ensue. The fundamental environmental condition to which Hull alludes regards the existence of laws, which, acting as trans-contextual constraints on belief-formation, would be sufficient, in his view, to force convergence and eliminate our socially and culturally ingrained biases. In brief, if there are laws then progress is possible. Of course, always in line with the hypothetical nature of EET, this is a hypothetical argument.

Campbell assimilates scientific progress to the increased verisimilitude of our picture of reality. This is a view that other eliminativists have endorsed (Popper and Kitcher). In section 5.2.2 we saw that all BVSR processes are eliminative. However, the commitment to eliminativism does not guarantee the increased approximation to the truth of our hypotheses, as repeatedly shown by critics of Popper (e.g. Howson 2001 ch. 5) and Kitcher (e.g. Rosenberg 1996). There is no assurance that the eliminative process can be proved to be deductively valid in science since not all possible alternative hypotheses compatible with the given evidence can be generated and thus eliminated. Even if we lower our expectations by abandoning the quest for a deductively valid procedure I remain sceptical about the chances of eliminativism to demonstrably lead to increased verisimilitude. This would be possible if we had good reasons to believe that all plausible hypotheses have been proposed, fully formulated and seriously evaluated. Unfortunately we know that in practice hypotheses are selected against for bad reasons and that potentially interesting hypotheses are not fully articulated. Science is an historical process, and like any historical process is affected by contingent factors. Such contingency is a sufficient reason not to be optimistic about the possibility for humans to generate all plausible hypotheses given a potentially constantly changing amount of evidence. With this I am not retracting what I argued for in sections 5.2.2 and 5.2.3 (i.e. that eliminate inductively is the best we can do in order to gain knowledge), but simply pointing out that the patchy eliminativism of actual scientific practice does not guarantee that science is approximating the correct picture of reality.

In my opinion, evolutionary and more generally naturalistic considerations suggest that a minimal realist approach is more consonant to EET. Axiologically, instrumentalism better suits the EET stance than realism for the following reason. Naturalism treats axiology as part of the theory of inquiry. The aim of axiology is to assess what cognitive aims are achievable and which are not. Instrumental reliability is clearly an achievable aim while truth is not clearly so. A fortiori, reliability is a legitimate aim of inquiry.

As far as scientific realism is concerned, EET should find solace in the instrumental reliability of science and remain agnostic about more global scenarios. Agnosticism could also be praised because it arguably provides a better reply than realism to the relativist. In fact, while a realist must show that among the non-formulated *possible* alternatives none is as empirically adequate as the actual one, the agnostic can shift the burden of proof and challenge the relativist to find a better one (Rosenberg 1996).

My conclusion is that EET can show that science is progressive in the sense of improving its ability to produce reliable knowledge. EET can explain why this is so by reference to the standards of evaluation that EET identifies (cf. section 5.3). These standards are not negotiable but in principle revisable. The existence of such standards coupled with the thesis of ontological realism is sufficient, I believe, to dispel any further doubt concerning the relativism of my view.

But even though we can show that science somehow progresses, I remain agnostic about the quality of such progress. It remains an open question whether our knowledge is globally progressing, and whether our present picture of reality will converge on the real description as science evolves. This latter realist view should be considered distinct from the genuine and sensible hypothetical realism advocated by EET.

Conclusion

In the previous pages I have tried to assess the virtues and shortcomings of the EET perspective to epistemology. As far as the virtues are concerned, three themes have emerged. First, I believe to have shown that EET models cannot be dismissed on a priori grounds just because scientific evolution is a different selective process from biological evolution. Scientific and biological evolution are partly disanalogous processes, but this fact is of no epistemological relevance. Secondly, I have assessed whether the evolutionary approach to epistemology proposed by EET contributes to answering the challenges typical of epistemic naturalism. In this respect I believe to have shown that EET's approach is genuinely normative and not seriously relativistic. Thirdly, in the thesis I have also tried to defend the general validity of the blind variation and selective retention model, which seems to me at the basis of any genuine evolutionary approach to epistemology.

However, I have also argued in favour of the complementarity and peaceful co-existence between EET and some more traditional approaches. This epistemological pluralism is an outcome of the partial incompleteness of EET I have been outlining.

My thesis will have reached its aim if the reader realises that the adventure Campbell started (I have used his EET model as my "paradigm"), is worth exploring, if the reader only temporarily marvels at the novel and original insights that an evolutionary epistemological perspective offers. What is certainly true is that the study of science and epistemology cannot be completely sought by means of a purely abstract, disembodied, logical and individualistic perspective. EET models offer original means to obviate some of the limits of these approaches. Even though the value of EET will ultimately depend on both the vindication of the BVSR model, and on its application to concrete cases, I believe that the scientific study of science has received a major boost with the advent and emergence of EET's adaptationist programme.

Bibliography

1) Allen, G.E. (1991) - *Essay Review of Hull's "Science as a Process"* – ISIS, 82, pp.698-704.

2) Amundson, R. (1989) – *The Trials and Tribulations of Selectionist Explanations* – pp. 413-432 in Hahlweg, K. & Hooker, C.A. eds. (1989)

3) Ayala, F.J. & Dobzhansky, T. eds (1974) - *Studies in the Philosophy of Biology* - MacMillan, London

4) Barnes, B. & Bloor, D. (1982) – *Relativism, Rationalism and the Sociology of Knowledge* – pp. 21-47 in Hollis, M. & Lukes, S. eds. (1982) "Rationality and Relativism", Oxford: Basil Blackwell.

5) Boyd, R. & Richerson, P. (1985) - *Culture and the Evolutionary Process* – University of Chicago Press

6) Bradie, M. (1986) - *Assessing Evolutionary Epistemology* – Biology and Philosophy, 1, pp. 401-459

7) Bradie, M. (1989) - *Evolutionary Epistemology as Naturalized Epistemology* – pp. 393-412 in Hahlweg, K. & Hooker, C.A. eds. (1989)

8) Bradie, M. (1990) - *The Evolution of Scientific Lineages* – PSA 1990, vol.2, p.245-254

9) Bradie,M. (2001) - *The Metaphysical Foundation of Campbell's Selectionist Epistemology* - pp. 35-52 in Heyes, C. & Hull, D. eds. (2001)

10) Campbell,D.T. (1974a) - *Evolutionary Epistemology* – pp. 413-463 in Schilpp, P.A. ed. (1974)

11) Campbell, D.T. (1974b) - *Unjustified Variation and Selective Retention in Scientific Discovery* – pp. 139-161 in Ayala, F.J. & Dobzhansky, T. eds. (1974)

12) Campbell, D.T. (1974c) - *Downward Causation* – pp. 179-186 in Ayala, F.J. & Dobzhansky, T. eds. (1974)

13) Campbell, D.T. (1977) - *Descriptive Epistemology: Psychological, Sociological and Evolutionary* – chapter 17 in Campbell, D.T. (1988) - "Methodology and epistemology for social science: Selected papers", University of Chicago Press

14) Campbell, D.T. (1979) - *A Tribal Model of the Social System Vehicle that Carries Scientific Knowledge* - chapter 18 in Campbell, D.T. (1988) - "Methodology and epistemology for social science: Selected papers", University of Chicago Press

15) Campbell, D.T. (1981) - *Science's Social System of Validity Enhancing Collective Belief Change* - chapter 19 in Campbell, D.T. (1988) - "Methodology and epistemology for social science: Selected papers", University of Chicago Press

16) Campbell,D.T. (1994)- *How Individual and Face-to-Face-Group-Selection undermine Firm Selection in Organizational Evolution* – pp. 23-38 in Baum, J.A.C. & Singh, J.V. eds. (1994) "Evolutionary Dynamics of Organization", Oxford University Press

17) Campbell, D.T. and Paller B.T. (1989) - *Extending evolutionary epistemology to 'justifying' scientific beliefs* – ch. 5 in Hahlweg, K. & Hooker, C.A. eds. (1989)

18) Campbell, D.T. (1997) – *From Evolutionary Epistemology via Selection Theory to a Sociology of Scienfitic Validity* – Evolution and Cognition, 3, pp.5-38

19) Carey, T.V. (1998) - *The Invisible Hand of Natural Selection, and Vice Versa* - Biology and Philosophy, 13, pp.427-442

20) Carnap, R. (1936) - *Testability and Meaning* – Philosophy of Science, 3, pp. 419-471

21) Cavalli-Sforza, L. & Feldman, M. (1981) - *Cultural transmission and evolution: a quantitative approach* – Princeton University Press

22) Chomsky, N. (1987) – *Language and Problems of Knowledge: The Managua Lectures* – MIT Press

23) Cohen, L.J. (1973) - *Is the progress of science evolutionary?* – British Journal for the Philosophy of Science, 24, pp. 41-61

24) Cosmides, L. & Tooby, J. (1994) - *Better than Rational: Evolutionary Psychology and the Invisible Hand* - AEA Papers and Proceedings, pp. 327-332
25) Cziko, G. (1995) - *Without miracles: Universal Selection Theory and the Second Darwinian Revolution* – MIT Press
26) Cziko, G. (2001) - *Universal Selection Theory and the Complementarity of Different Types of Blind Variation and Selective Retention* – pp. 15-35 in Heyes, C. & Hull, D. eds. (2001)
27) Dawkins, R. (1980) - *Universal Darwinism* – pp. 403-425 in Bendall, D.S. ed. (1980) "Evolution from Molecules to Men", Cambridge University Press
28) Dawkins, R. (1989) - *The Selfish Gene* – Oxford University Press
29) Dawkins, R. (1999) - *The Extended Phenotype* – Oxford University Press
30) Dennett, D.C. (1982) - *The Intentional Stance* – MIT Press
31) Dennett, D.C. (1991) - *Consciousness Explained* – Boston, Little Brown
32) Dennett, D.C. (1996) - *Darwin's Dangerous Idea: Evolution and the Meanings of Life* - Penguin Books
33) Dennett, D.C. (2002) - *The Baldwin Effect: a Crane, not a Skyhook* – in Weber, B. & Depew, D. eds. (2002) "Learning and Evolution: the Baldwin Effect Reconsidered", MIT Press
34) Diamond, J. (1998) - *Guns, Germs and Steel* - Vintage
35) Dobzhansky, T. (1974) - *Chance and creativity in evolution* – in Ayala, F.J. & Dobzhansky, T. eds. (1974)
36) Donoghue, M.J. (1990) - *Sociology, Selection and Success* - Biology and Philosophy, 5, pp. 459-72
37) Downes, S.M. (1999) - *Can Scientific Development and Children's Cognitive Development Be the Same Process?* – Philosophy of Science, 66, pp.565-578
38) Downes, S.M. (2002) - *Baldwin Effects and the Expansion of the Evolutionary Repertoire of Evolutionary Biology* - in Weber, B. & Depew, D. eds. (2002) "Learning and Evolution: the Baldwin Effect Reconsidered", MIT Press

39) Dretske, F.I. (1971) - *Perception from an Epistemological point of view* – The Journal of Philosophy, 66, pp. 584-591

40) John Dupre' (1990) - *Scientific Pluralism and the Plurality of the Sciences* - Philosophical studies, 60, pp. 61-76

41) Dupre', J. (1993) – *The Disorder of Things: Metaphysical Foundations of the Disunity of Science* – Harvard University Press

42) Dupre', J. (2001) - *Human Nature and the Limits of Science* – Oxford University Press

43) Feyerabend, P. (1975) - *How to defend Society against Science* – Radical Philosophy, 11, Summer 1975

44) Giere, R.N. (1985) - *Philosophy of Science Naturalized* - Philosophy of Science, 52, pp. 331-356

45) Giere, R.N. (1996) - *The Scientist as Adult* – Philosophy of Science, 63, pp. 538-541

46) Giere, R.N. (2000) - *Naturalized Philosophy of Science* - entry in The Routledge Encyclopaedia of Philosophy

47) Giere, R.N. (2001) - *Critical Hypothetical Evolutionary Naturalism* - pp.53-70 in Heyes, C. & Hull, D. eds. (2001)

48) Godfrey-Smith, P. (2002) - *Between Baldwin Skepticism and Baldwin Boosterism* – in Weber, B. & Depew, D. eds. (2002) "Learning and Evolution: the Baldwin Effect Reconsidered", MIT Press

49) Goldman, A. (1986) - *Epistemology and Cognition* - Harvard University Press

50) Gopnik, A. (1996) - *The Scientists as Child* – Philosophy of Science, 63, pp. 485-514

51) Griffiths, P.E. (2002) - *What is Innateness?* – The Monist, 85, pp.70-85

52) Gould, S.J. (1980) - *Is a New and General Theory of Evolution Emerging?* – Paleobiology, 6, pp. 119-130

53) Gould, S.J. & Lewontin, R.C. (1979) - *The Spandrels of San Marco and the Panglossian Paradigm: a Critique of the Adaptationist Programme* – in Sober, E. ed. (1993) "Conceptual Issues in the Philosophy of Biology", MIT Press

54) Hahlweg, K. & Hooker, C.A. eds. (1989) - *Issues in Evolutionary Epistemology* - State University of New York Press

55) Hempel, C.G. (1965) - *The Theoretician's Dilemma* – in Hempel, C.G. (1965) "Aspects of Scientific Explanation", New York Free Press

56) Heyes, C. & Hull, D. eds (2001) - *Selection theory and social construction: the evolutionary naturalistic epistemology of Donald T. Campbell* – SUNY Press

57) Howson, C. (2000) - *Hume's problem: Induction and the Justification of Belief* - Oxford University Press

58) Hull, D.L. (1982) - *The Naked Meme* – pp. 273-327 in Plotkin H.C. ed. (1982) "Learning, Development and Culture", New York, John Wiley & Sons

59) Hull, D.L. (1986) - *On Human Nature* – in Ruse, M. & Hull, D.L. ed.s (1998) "Philosophy of Biology", Oxford University Press

60) Hull, D.L. (1988a) - *Science as a Process* – The University of Chicago Press

61) Hull, D.L. (1988b) - *A Mechanism and Its Metaphysics* - Biology and Philosohy, 3, pp.123-155

62) Hull, D.L. (1990) - *Conceptual Selection* – Philosophical Studies, 60, pp.77-87

63) Hull, D.L. (1992a) - *Testing Philosophical Claims about Science* – PSA 1992, vol.2, pp.468-75

64) Hull, D.L. (1992b) - *The Evolution of Conceptual Systems in Science* – World Futures, 34, pp. 67-82

65) Hull, D.L. (1994) - *That just don't sound right: a plea for real examples in philosophy of science* – ch. 10 in Hull, D.L. (2001)

66) Hull, D.L. (1996) - *Why Scientists Behave Scientifically* – MRS Bulletin, p.72 ff

67) Hull, D.L. (1997) - *What's Wrong with Invisible-Hand Explanations?* – Philosophy of Science (Proceedings), 64, pp. 117-126

68) Hull, D.L. (1999)- *The Use and Abuse of Sir Karl Popper* – Biology and Philosophy, 14, pp. 481-504
69) Hull, D.L. (2001a) - *Interactors versus Vehicles* – ch. 1 in Hull, D.L. (2001d)
70) Hull, D.L. (2001b) - *Taking Vehicles Seriously* – ch. 2 in Hull, D.L. (2001d)
71) Hull, D.L., R.Langman, R. & Glenn, S. (2001c) - *A General Account of Selection: Biology, Immunology and Behaviour* – ch. 3 in Hull, D.L. (2001d)
72) Hull, D.L. (2001d) - *Science and Selection: Essays on Biological Evolution and the Philosophy of Science* – Cambridge University Press
73) Hull, D.L. (2001e) – *In Search of Epistemological Warrant* – ch. 9 in Heyes, C. & Hull, D. eds (2001)
74) Kary, C.E. (1982) - *Can Darwinian Inheritance be extended from Biology to Epistemology?* - PSA, Vol. 1982, pp.356-369
75) Kim, J. (1988) - *What is 'Naturalized Epistemology'?* - Philosophical Perspectives, 2, Epistemology, pp. 381 -405.
76) Kimura, M. (1983) - *The Neutral Theory of Molecular Evolution* - Cambridge University Press
77) Kitcher, P. (1988) - *Selection Among the Systematists* – Nature, 336, pp. 277-8
78) Kitcher, P. (1992) - *The Naturalists Return* – Philosophical Review, 101, pp. 53-114
79) Kitcher, P. (1993) - *The Advancement of Science. Science without Legend, Objectivity without Illusions* – Oxford University Press
80) Kuhn, T.S. (1970a) – *The Structure of Scientific Revolutions* – The University of Chicago Press
81) Kuhn, T.S. (1970b) - *Logic of Discovery or Psychology of Research?* – pp. 1-23 in Lakatos, I. & Musgrave, A., eds. (1970) *Criticism and the Growth of Knowledge,* Cambridge University Press.

82) Kuhn, T.S. (1970c) - *Reflections on My Critics* –pp. 231-278 in Lakatos, I. & Musgrave, A., eds. (1970) *Criticism and the Growth of Knowledge,* Cambridge University Press.

83) Kuhn, T.S. (1977) - *Objectivity, Value Judgments and Theory Choice* – pp. 320-339 in Kuhn, T.S. "The Essential Tension", The University of Chicago Press

84) Lakatos, I. (1970) - *Falsificationism and the Methodology of Scientific Research programmes* – pp. 91 ff. in Lakatos, I. & Musgrave, A., eds. (1970) Criticism and the Growth of Knowledge, Cambridge University Press.

85) Lakatos, I. & Musgrave, A., eds. (1970) - *Criticism and the Growth of Knowledge* - Cambridge University Press.

86) Latour, B. (1987) - *Science in action* - Harvard University Press

87) Latour, B. (1990) - *Review of D. Hull's 'Science as a process"* - Contemporary Sociology, 19, pp. 281-2

88) Laudan, L. (1977) - *Progress and its Problems* – University of California Press

89) Laudan, L. (1980) – *Why was the Logic of Discovery abandoned?* - pp. 173-183 in Nickles, T. ed. (1980) "Scientific Discovery, Logic and Rationality", Reidel Dordrect

90) Laudan, L. (1987a) - *Progress or Rationality? The Prospects of Normative Naturalism* -American Philosophical Quarterly, 24, pp. 19-33

91) Laudan, L. (1987b) - *Relativism, Naturalism and Reticulation* - Synthese, 71, pp.221-234

92) Laudan, L. (1989) - *If It Ain't Broke, Don't Fix It* - British Journal for the Philosophy of Science, 40, pp. 369-375

93) Laudan, L. (1996) - *Beyond Positivism and Relativism* – Boulder, Westview Press

94) Leakey, R. (1994) - *The Origin of Humankind* – Weidenfeld and Nicolson

95) Lewontin, R.C. (1982) - *Organism and Environment* – pp.151-169 in H.C. Plotkin, ed. (1982) "Learning, Development and Culture", John Wiley & Sons Ltd.

96) Maffie, J. (1990a) - *Naturalism and the Normativity of Epistemology* –

Philosophical Studies, 59, pp. 333-349

97) Maffie, J. (1990b) - *Recent Work on Naturalized Epistemology* - American Philosophical Quarterly, 27, pp. 281-294

98) Magnus, D. (1998) - *Evolution Without Change in Gene Frequencies* – Biology and Philosophy, 13, pp. 255-261

99) Maynard-Smith, J. (1988) - *Mechanism of Advance: Book Review of Hull's Science as a Process* – Science, 242, pp.1182-3

100) Mayr, E. (1976) - *Evolution and the Diversity of Life* – Harvard University Press

101) Monod, J. (1971) - *Chance and Necessity* – Knopf

102) Neander, K. (1988) - *What does natural selection explain? Correction to Sober* - Philosophy of Science, 55, pp.422-426

103) Nozick, R. (1994) - *Invisible-Hand Explanations* - American Economic Review, May 1994, pp.314-18

104) Oldroyd, D. (1990)- *David Hull's Evolutionary Model for the Progress and Process of Science* - Biology and Philosophy, 5, pp.475-87

105) Pinch, T.J. (1989) - *Essay Review: Science as a Process* - Annals of Science, 46, pp.521-26

106) Plotkin, H.C. ed. (1982) – *Learning, development and Culture: Essays in Evolutionary Epistemology* – New York, John Wiley & Sons

107) Popper, K.R. (1963) - *Conjectures and Refutations: the Growth of Scientific Knowledge* - Routledge

108) Popper, K.R. (1966) - *The Open Society and Its Enemies* – Routledge

109) Popper, K.R. (1974a) - *Autobiography* – pp. 1-181 in Schilpp, P.A. ed. (1974)

110) Popper, K.R. (1974b) - *Darwinism as a Metaphysical Research Program* – pp. 133-143 in Schilpp, P.A. ed. (1974)

111) Popper, K.R. (1974c) - *Replies to my Critics* – pp. 1048-ff. in Schilpp, P.A. ed. (1974)

112) Popper, K.R. (1978) - *Natural Selection and the Emergence of Mind* – Dialectica, 32, pp. 339-355

113) Popper, K.R. (1979) - *Objective Knowledge: an Evolutionary Approach* – Oxford University Press

114) Putnam, H. (1975) - *Philosophy and our mental life* – ch. 14 in Putnam, H. (1975) "Mind, Language and Reality", Cambridge University Press

115) Quine, W.V.O. - *Epistemology Naturalized* – in Quine, W.V.O. (1969) "Ontological Relativity and Other Essays", Columbia University Press

116) Reichenbach, H. (1938) - *Experience and Prediction* – University of Chicago Press

117)Richards,R.J.(1981)- *Natural selection and Other Models in the Historiography of Science* – pp. 37-76 in M.B. Brewer, M.B. & Collins, B.E. eds. (1981) "Scientific Inquiry and the Social Sciences: A Volume in Honor of Donald T. Campbell", San Francisco, Jossey-Bass

118) Richards, R.J. (1987) - *Darwin and the Emergence of evolutionary Theories of Mind and Behaviour* – University of Chicago Press

119) Rorty, R. (1980) - *Philosophy and the Mirror of Nature* – Princeton University Press

120) Rosenberg, A. (1990) - *Normative Naturalism and the Role of Philosophy* - Philosophy of Science, 57, pp. 34-43

121) Rosenberg, A. (1992) - *Selection and Science: Critical Notice of David Hull's Science as a Process* - Biology and Philosophy, 7, pp. 217-228

122) Rosenberg, A. (1996) - *A Field Guide to Recent Species of Naturalism* - British Journal for the Philosophy of Science, 47, pp. 1-29

123) Ruse, M. (1977) - *Karl Popper's Philosophy of Biology* – Philosophy of Science , 44, pp. 638-661

124) Ruse, M. (1986) - *Taking Darwin Seriously: A Naturalistic Approach to Philosophy* - Oxford, Blackwell

125) Ruse, M. (1989) - *The View from Somewhere: a Critical Defense of Evolutionary Epistemology* – pp. 185-228 in Hahlweg, K. & Hooker, C.A. eds. (1989)

126) M. Ruse, M. (2001) - *On Being a Philosophical Naturalist: a Tribute to D.T. Campbell* – ch. 5 in Heyes, C. & Hull, D. eds. (2001)

127) Salmon, W.C. (1990) - *Rationality and Objectivity in Science, or Tom Kuhn meets Tom Bayes* – pp. 175-204 in Minnesota Studies in the Philosophy of Science, vol. XIV, University of Minnesota Press

128) Schilpp, P.A. ed. (1974) - *The Philosophy of Karl Popper* – La Salle, Open Court

129) Siegel, H. (1980) - *Justification, Discovery, and the Naturalizing of Epistemology* – Philosophy of Science, 47, pp.297-321

130) Siegel, H. (1990) - *Laudan's normative naturalism* - Studies in History and Philosophy of Science, 21, pp. 295-313.

131) Skolimowsky, H. (1974) - *Popper and Objectivity in Science* – pp. 483- 508 in Schilpp, P.A. ed. (1974)

132) Sober, E. (1978) – *Psychologism* – Journal for the Theory of Social Behaviour, 8, pp. 165-191.

133) Sober, E. (1980) - *Evolution, Population Thinking and Essentialism* – Philosophy of Science, 47, pp. 350-383

134) Sober, E. (1981) - *The Evolution of Rationality* – Synthese, 46, pp. 95-120

135) Sober, E. (1993) - *The Nature of Selection* – The University of Chicago Press

136) Sober, E. (1994a) - *The adaptive advantage of learning and a priori prejudice* – pp. 50-70 in Sober, E. (1994) "From a Biological Point of View", Cambridge University Press

137) Sober, E. (1994b) - *Prospects for an Evolutionary Ethics* – in Sober, E. (1994) "From a Biological Point of View", Cambridge University Press

138) Sober, E. (1994c) - *Models of Cultural Evolution* – in Sober, E. ed. (1994) "Conceptual Issues in Evolutionary Biology", second edition, MIT Press

139) Sober, E. (1994d) - *Progress and Direction in Evolution* – pp. 19-33 in Campbell,

J. ed. (1994) "*Creative Evolution?!* ", Jones and Bartlett

140) Sober, E. (1999a) – *Physicalism from a Probabilistic Point of View* – Philosophical Studies, 95, pp.135-174

141) Sober, E. (1999b) – *Multiple Realizability Argument against Reductionism* – Philosophy of Science, 66, pp. 542-564

142) Sober, E. (1999c) – *Testability* – Proceedings and Addresses of the APA, 73, pp. 47-76

143) Sober, E. (2000) - *Philosophy of Biology* – Westview Press

144) Sober, E. & Wilson, D.S. (1998) - *Unto Others: The Evolution and Psychology of Unselfish Behavior* – Harvard University Press

145) Sober, E., Wright. E. & Levine, A. (1992) – *Reconstructing Marxism: Essays on Explanation and the Theory of History* – Verso Books

146) Sterelny, K. (1995) - *Review of "The Adapted Mind"* by J. Tooby & L. Cosmides – Biology and Philosophy, 10, pp. 265-280

147) Stephens, C. (2001) - *When is it Selectively Advantageous to Have True Beliefs? Sandwiching the Better Safe than Sorry Argument* – Philosophical Studies, 105, pp. 161-189

148) Stich, S.P. (1985) - *Could Man be an Irrational Animal?* – Synthese, 64, pp. 115-135

149) Stich, S.P. (1990) - *The Fragmentation of Reason* – MIT Press

150) Stich, S.P. (2000) - *Epistemic Relativism* – entry in The Routledge Encyclopaedia of Philosophy

151) Stokes, G. (1989) - *From Physics to Biology: Rationality in Popper's Conception of Evolutionary Epistemology* – in Hahlweg, K. & Hooker, C.A. eds. (1989)

152) Suppe, F. (1974) - *The Structure of Scientific Theories* – University of Illinois Press

153) Tomasello, M. (1997) - *Two hypotheses about primate cognition* – ch. 9 in Tomasello, M. & Call, J. (1997) "Primate Cognition", Oxford University Press

154) Uebel, T.E. (1991) - *Neurath's Programme for Naturalistic Epistemology* – Studies in the History and Philosophy of Science, 22, pp. 623-646

155) Uebel, T.E. (1996) - *Anti-Foundationalism and the Vienna Circle's Revolution in Philosophy* – British Journal for the Philosophy of Science, 47, pp. 415-440

156) Vining, D.R. Jr. (1986) - *Social versus Reproductive Success: The Central Theoretical Problem of Human Sociobiology* – The Behavioral and Brain Sciences, 9, pp. 167-216

157) Williams, G.C. (1982) - *A Comment on Genetic Assimilation* – ch. 11 in Plotkin, H.C. (ed.) (1982)

158) Wilson, D.S. (1989) – *Levels of Selection: An Alternative to Individualism in Biology and the Human Sciences* – Social Networks, 11, pp. 257-272

159) D. S. Wilson, E. Dietrich & A. B. Clark (2003) – *On the inappropriate use of the naturalistic fallacy in evolutionary psychology* – Biology and Philosophy, 18, pp. 669–682

160) Wilson, E.O. and Lumsden, C.J. (1981) - *Genes, Mind and Culture: the Coevolutionary Process* – Harvard University Press

161) Worrall, J. (1990) - *Rationality, Sociology and the Symmetry Thesis* - International Studies in the Philosophy of Science, 4

162) Worrall, J. (1989) - *Fix it and be damned: a reply to Laudan* - British Journal for the Philosophy of Science, 40, pp. 376-388

163) Worrall, J. (1999) – *Two cheers for Naturalised Philosophy of Science: or why Naturalised Philosophy of Science is not the Cat's Whiskers* – Science and Education, 8, pp. 339-361.

Acknowledgments

I wish to thank sincerely my supervisor, Carl Hoefer, for his extreme patience and intellectual contribution. I also need to thank my co-supervisor, Elliott Sober, for his precious comments, David Hull for his helpful advice, and John Worrall for trying to convince me to give up naturalism.

I dedicate this thesis to Sangeetha, Ramona Elsa, and my two families.