A Mechanism and Its Metaphysics: An Evolutionary Account of the Social and Conceptual Development of Science*

DAVID L. HULL

Department of Philosophy
Northwestern University
Evanston IL 60201

ABSTRACT: The claim that conceptual systems change is a platitude. That our conceptual systems are theory-laden is no less platitudinous. Given evolutionary theory, biologists are led to divide up the living world into genes, organisms, species, etc. in a particular way. No theory-neutral individuation of individuals or partitioning of these individuals into natural kinds is possible. Parallel observations should hold for philosophical theories about scientific theories. In this paper I summarize a theory of scientific change which I set out in considerable detail in a book that I shall publish in the near future. Just as few scientists were willing to entertain the view that species evolve in the absence of a mechanism capable of explaining this change, so philosophers should be just as reticent about accepting a parallel view of conceptual systems in science "evolving" in the absence of a mechanism to explain this evolution. In this paper I set out such a mechanism. One reason that this task has seemed so formidable in the past is that we have all construed conceptual systems inappropriately. If we are to understand the evolution of conceptual systems in science, we must interpret them as forming lineages related by descent. In my theory, the notion of a "family resemblance" is taken literally, not metaphorically. In my book, I set out data to show that the mechanism which I propose is actually operative. In this paper, such data is assumed.

KEY WORDS: conceptual evolution, scientific theories, selection processes, social organization of science.

> *Scientists want credit for their contributions? My dear, let us hope that it isn't true! But if it is true, let us hope that it doesn't become widely known.*

The idea of treating conceptual change as resulting from a selection process has seemed intriguing ever since Darwin established that biological species evolve primarily by means of natural selection. However, not until recently has anyone presented anything like a detailed application of the basic models of population genetics to conceptual change (Alexander 1979, Cavalli-Sforza and Feldman 1981, Lumsden and Wilson 1981, Schilcher and Tennant 1984, Boyd and Richerson 1985). Some of the

Biology and Philosophy **3** (1988) 123–155.

models are literal applications. Various behaviors are treated as being genetically influenced phenotypic characters. Because my concern is conceptual change in science, literal applications are not liable to get me very far. We are a curious species. Our chief adaptation is playing the knowledge game. Our propensity to learn about the world in which we live is surely based in our genetic make-up. We might even be programmed to perceive the world in certain ways. For example, apparently people tend to divide up the visible spectrum along very similar lines, regardless of the contingencies of their culture (Berlin and Kay 1969, Bornstein 1973). Viewing the world essentialistically may also be part of our "hard wiring." If we are not born essentialists, at the very least we come to adopt this perspective quite rapidly.

We are also a social species. Our preference for living in herds is very likely to have some genetic basis. Since science itself is a social institution with its own norms and organization, our tendency to form groups also contributes to it. Even our general ability to use language must have some genetic basis, and language is yet another prerequisite for science. 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 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 (for the literature on evolutionary epistemology, see Campbell 1974a, Bradie 1986, Campbell et al. 1987, Plotkin 1987).

In the past most authors who have treated cultural evolution in general and scientific change in particular as selection processes have taken gene-based natural selection as the exemplar and reasoned analogically to conceptual change. However, a more appropriate strategy to to present a general anslysis of selection processes equally applicable to all sorts of selection processes (Toulmin 1972, Campbell 1974a). After all, the reaction of the immune system to antigens is an instance of a selection process which differs just as radically from gene-based natural selection as does conceptual change in science. Any analysis of selection processes must apply to it as well as to natural selection. Such a general analysis must be sufficiently general so that it is not biased toward any particular sort of selection process but not so general that any and all natural processes turn out to count as instances of selection. Biological evolution, the reaction of the immune system to antigens, and cultural learning must count but not lead balls rolling down inclined planes or the planets circling the sun.

Even if such a general analysis of social learning as a selection process is successful, it alone is not enough to account for conceptual change in science. The most puzzling feature of science is that it works so well in realizing its manifest goals, so much better than any other social institution. By and large scientists really do what they claim to do. All social institutions exhibit norms, but even when these social norms are translated from their usual hypocritical formulations to accord more closely with the norms that are actually operative, individual infractions are common. In science they are correspondingly rare. In 1985 at a meeting of the American Association for the Advancement of Science, William F. Raub, Deputy Director for Extramural Research and Teaching at the National Institute of Health, presented a paper in which he reported detecting only fifty cases of serious fraud out of 20,000 projects supported by the institute during a five year period. Most of these cases involved such things as sloppy record keeping, the rest outright falsification of data. Observations in science are certainly, to some extent, theory laden, but the experiments which scientists run are not just *pro forma* rituals designed to support preconceived conclusions. By and large, scientists really do run the experiments that they claim to run, present their findings accurately enough, and are influenced by the results. They even acknowledge the contributions which other scientists make (for additional estimates of the amount of various sorts of "fraud" in science, see, Zuckerman 1977, Culliton 1983, 1986, 1987, Norman 1984, 1987, Marshall 1986, Koshland 1987, and White 1987).

Although we are currently experiencing a period of science bashing, even the most hysterical critics have been unable to come up with rates of "fraud" in science that begin to come close to those repeatedly documented in other social institutions. Either scientists are not guilty of the substantial amounts of fraud common in other professions or else they are extremely good at covering it up. The romantic view of scientists as dispassionate, disinterested seekers after truth for its own sake needs debunking, but not because science is a plot against "The People," not simply because it is overly idealized, but because there is some danger, though not much, that scientists might actually be tempted to put this romantic view into practice. If scientists at large adhered to the professed mores of science, science might be possible but I doubt it. As it is, scientists have too much sense to behave too strictly in accordance with the way that officially they are supposed to behave.

Numerous authors before Darwin suggested that species might evolve. A few even suggested mechanisms for this evolution. Darwin was so much more efficacious than his predecessors for a variety of reasons. He not only suggested a mechanism to account for the transformation of species but also worked it out in great detail. It is one thing to sketch the general outlines of natural selection, as both Wallace and Darwin did in their Linnean Society papers. It is quite another to explain how the extraordinary characteristics of sterile castes in eusocial organisms might have resulted from natural selection as well as dozens of other problem cases. Darwin was especially adept at transforming anomalies into confirming instances.

In addition, Darwin presented massive amounts of data to support his theory, much of it obtained from the work of others, and set out his views in a scientifically respectable way. The *Origin of Species* makes for slow reading because it contains so many examples, so much data, but Darwin did not stop with the *Origin*. He spent the next two decades publishing book after book in which he investigated one aspect of evolution after another, each laden with masses of data. Darwin's theory might be wrong, but no one could claim on Darwin's death that he did not take it seriously. More than this he intertwined his professional career with those of other scientists, scientists who stood to gain or lose depending on the fate of Darwin's revolutionary theory. Strangely enough, Darwin's presentation of natural selection as the primary mechanism for the transmutation of species legitimized the idea of organic evolution even though most of his contemporaries rejected it as the chief mechanism for evolutionary change (Ghiselin 1969, Ruse 1979).

The standards for scientific acceptability have hardly decreased since Darwin's day. Anyone who expects an evolutionary account of conceptual development in science to be taken seriously must attempt to meet these standards. In this short paper, however, all I can present is a sketch, along with the promissory note that in my book (Hull 1988a) I set out a more detailed account along with the necessary data.[1] Instead of explaining such things as the slave making habits of ants, I have to explain why scientists on occasion behave so altruistically, for instance, by giving credit to their closest competitors, and why on occasion they do not. Under what conditions are scientists likely to give credit to other scientists, and under what conditions are they likely to claim priority for themselves? As one might imagine, I do not think that "intellectual justice" is the entire story. Scientist do strive to behave the way that they "should," but other factors are operative as well, in fact more operative.

Much of what I have to say about the mechanisms that drive science has been said before. The two innovations which I introduce to explain the behavior of scientists are *conceptual inclusive fitness* and the *demic structure of science*. Just as organisms behave in ways which result in replicates of their own genes or duplicates of these genes in close kin being transmitted to later generations, scientists behave in ways calculated to get their views accepted as their view by other scientists, in particular those scientists working on problems most closely associated with their own. Scientists also tend to organize themselves into tightly knit, though relatively ephermeral research groups to develop and disseminate a particular set of views. On the account that I am urging, conceptual change in science should be most rapid when scientists are subdivided into competing research groups. The factionalism that scientists themselves so often decry facilitates rather than frustrates progress in science (Blau 1978).

Science is inherently both a competitive and cooperative affair. If all that mattered was a scientist's individual conceptual inclusive fitness, the relationships between scientists would be complicated enough, but when alliances and allegiances are added, the story gets even more complicated. However, these complications cannot be ignored because science *is* essentially social. Individuals can learn about the world in which they live by confronting it directly, but if science is to be cumulative, social transmission is necessary. In addition, the sort of objectivity that gives science its peculiar character is a property of social groups, not isolated investigators.

CONCEPTUAL INCLUSIVE FITNESS

Science is a matter of both cooperation and competition. Both sorts of behavior are natural; both need explaining. How come scientists can cooperate as much as they do in such a competitive activity? The most important sort of cooperation that occurs in science is the use by one scientist of the results of the research of other scientists. The worst thing that one scientist can do to another is to ignore his or her work; the best thing is to incorporate that work into one's own with an explicit acknowledgement. Somewhere in between is use without explicit acknowledgement. Scientists want their work to be acknowledged as original, but just as importantly, they want it to be accepted. To get it accepted, scientists must gain support from other scientists. One way to gain support is to show that one's own work rests solidly on preceding research, but the price of this support is a decrease in apparent originality. One cannot gain support from a particular work unless one cites it, and this citation automatically both confers worth on the work cited and detracts from one's own originality. Scientists would like total credit and massive support, but they cannot. Science is so organized that scientists are forced to trade off credit for support.

Several factors influence which way a scientist opts in particular circumstances. For example, scientists whose support is worth having are likely to be cited more frequently than those whose support is worth little or nothing. The preceding tendency is accentuated by the fact that there is little point in omitting reference to the contributions of a famous scientist because everyone who counts already knows who the author of this contribution is. When one fails to cite a well-known author, one gains neither credit nor support. The same observations do not hold for those lower in the scientific hierarchy, in particular one's graduate students. In the short term, their support is worth very little. However, graduate students are not entirely powerless because they are likely to be the chief conduits for one's work to later generations. Their success increases one's own conceptual inclusive fitness. Parent-offspring conflicts should be as common in science as elsewhere, and the resolutions of such conflicts should have the same general features.

Just as organisms in general behave in ways likely to increase their own genetic inclusive fitness, scientists tend to behave in ways calculated to increase their own conceptual inclusive fitness. In neither case are the entities involved necessarily aware of what they are doing. Flour beetles are unaware that such a thing as genetic inclusive fitness even exists let alone capable of performing the required calculations. Scientists are not appreciably different in their quest for conceptual inclusive fitness. The functioning of science had better not depend crucially on widespread self-awareness among scientists about their own motivations or the effects of their actions because most scientists are no more reflective about the ongoing process of science than are other professionals about their professions. Self-delusion rules.

So we are to believe, hospitals are run primarily for the good of patients, universities are run primarily for the good of students, and governments exist primarily to serve the people. Anyone who believes any of the preceding claims knows very little about hospitals, universities, or governments. Patients, students, and citizens do get some benefit from the relevant institutions, but when conflicts of interest arise, they are not always reconciled in the favor of those at the bottom of the hierarchy. In fact, they rarely are, self-serving rationalizations of the most embarrassing sort notwithstanding. However, I have found that scientists are a good deal less resistant to looking at their profession realistically than are members of other professions, primarily because even when the hypocrisy and romanticism are stripped away, science still retains its traditional characteristics. Science does have the characteristics traditionally attributed to it but not for the reasons usually given.

By definition, professions are occupations which are self-policing. One thing is clear about professions: by and large they do not police themselves very well, at least not according to the criteria that they profess. However, scientists police themselves as if science were made up of nothing but neighborhood yentas. In all social institutions both individual and group interests exist. They need not conflict, but all too often they do. We are all constantly being urged to behave in ways that are for the good of some group or other. Such appeals do have some effect, but too often not enough. University professors really should devote more time to undergraduate teaching. Those of us who are actively engaged in research periodically feel guilty about not spending more time on teaching and produce all sorts of justifications whose effect is to make us feel less guilty, but that is about all. Scientists adhere to the norms of science so well because, more often than not, it is of their own best self interest to do so. By and large, what is good for the individual scientist is actually good for the group. The best thing that a scientist can do for science as a whole is to strive to increase his or her own conceptual inclusive fitness.

Such striving is kept within certain bounds by two factors: the need of scientists to use each other's work and the possibility of empirical testing. However, in order to evaluate the manner in which these factors function in science, two ways in which a scientist can "sin" against the mores of science must be distinguished. The first is by publishing faulty work, whether intentionally or unintentionally ("lying"). The second and most common is the failure to give credit where credit is due ("stealing"). Scientists "lie" when they publish views which either they know to be false or else for which they have failed to perform the work necessary to warrant their making these views public. In everyday life, guilt is a function of intent. Getting in a car and intentionally running someone down is much worse then intentionally getting drunk and unintentionally running someone down, even though the person may be equally dead in both situations. Officially at least, a parallel distinction exists in science. Publishing fabricated data is the worst sin that a scientist can commit. Publishing the results of sloppy research is considered to be bad but somehow not quite so bad, even though the effects on anyone using fabricated and sloppy research are indistinguishable.

One explanation for this distinction in science is that science is not an activity that lends itself to rote procedures. Scientists cannot publish the truth and nothing but the truth because, at the cutting edge of science, no one can tell for sure what this truth is. To be successful in their investigations, scientists must be adept at judicious finagling. Time and again in the course of their investigations, scientists must exercise judgment. Sometimes these decisions are patent. If one spills coffee over a culture of mouth protozoa and they all die, no one would object to a scientist omitting this data point. But other decisions are not so obvious. Should one correct for the age of the culture? Perhaps young cultures are more (or less) resistant to stannous flouride than older cultures. Scientists try to account for or discount as many of these contingencies as they can, but time, money, energy, and insight limit this activity. The boundary between understandable error and inexcusably sloppy work is quite fuzzy. All scientists know that at times they have erred in this respect. Although those scientists who blatantly fabricate data are well aware of what they are doing, the boundary between error and fabrication can also become fuzzy. In certain circumstances, the only way to distinguish between the two would be extensive psychoanalysis. Why are Mendel's data so much better than they should be? Did he consciously fudge the results of his investigations or did he unconsciously classify borderline cases to enhance the relation between expected and observed values? All the efforts of the Mendel industry notwithstanding, we are unlikely ever to know.

Intent to one side, publishing work that other scientists use and find to be mistaken is punished severely in science, much more severely than "stealing," i.e., trying to pass off someone else's work as one's own. It is also, by all indications, much rarer. Why is lying so much rarer in science than stealing? Because it is punished so much more severely. Why is it punished so much more severely? Because stealing hurts only the person whose work has been appropriated, while lying hurts anyone who uses this work. Misassigned contributions are just as useful as work whose authorship is attributed correctly. However, the preceding comment concerns single instances. If the misassignment of credit in science were to become commonplace, the system would be seriously jeopardized.

The career of Sir Cyril Burt is a good case in point. When other scientists thought that all he had done was appropriate to himself the work done by his assistants, no one was especially excited. After all, that is what assistants are for. But when it began to appear that he had fabricated not only these assistants but also their research, his fellow scientists became more than a little anxious because it brought into doubt all of the work that they had published which was based on his fabricated results. Fabricated results need not be any more mistaken than the results of sloppy research, but there is no reason to expect them to be correct either. As science has been conducted in the past couple of centuries, by and large scientists have rewarded fairly consistently for doing what they are supposed to do and punished just as consistently when they transgress the actual mores of science.

As important as individual conceptual fitness is in science, conceptual demes also play an important role. Because few scientists have all the skills and knowledge necessary to solve the problems that they confront, they tend to band together to form research groups of varying degrees of cohesiveness. One function of these research groups is the sharing of conceptual resources (Giere 1988). These "demes" tend to be extremely ephemeral. They form and dissipate before anyone is liable to notice that they exist. However, every once in while, one will seem to have made some sort of headway or break through. A flurry of activity ensues, generally with little effect. Most activity in science, whether individual or group, has little discernible effect on science (for a critical review of this literature, see Fox 1983). However, on occasion, one of these groups is successful in the sense that others notice its achievements and either refute or adopt them. When the former occurs, interdemic selection is superimposed upon individual inclusive fitness. When the latter occurs, a set of views that originate in a small research group becomes widely disseminated, and interdemic selection is replaced by mass selection. As a result, only those research groups in science that are successful or significant failures are liable to be noticed.

Scientists do not simply read the literature to discover the truth. Rather they read it with an eye for work that bears on their own research. If a particular finding supports their own research, they are liable to incorporate it without testing. Testing is reserved for those findings which threaten one's own research. This tendency is accentuated in research groups. The scientists most likely not only to adopt one's views but also to be harmed most if they prove to be mistaken are one's allies. Some commentators on science urge scientists to test any and all work before they utilize it. They insist that a scientist should personally check everything that goes into any paper that bears his or her name (Broad and Wade 1982). I cannot conceive of worse advice. The whole point of scientists working together is to pool conceptual resources. Warm feelings to one side, cooperation in science is behaviorally indistinguishable from mutual exploitation. Perhaps one worker is mathematically quite adept but has very poor hands, while another can contribute very clean data sets even though he or she has to take the mathematics on faith. The best advice for a scientist who begins to doubt the reliability of the work of a colleague is to severe professional ties. One cannot waste one's time checking the result of one's co-workers.

Initially, criticism and evaluation come from within a research group. After publication it shifts to scientists outside the group, in particular to one's opponents. Individual scientists are to some extent objective. They know that, if they are lucky, their work will be held up to scrutiny. Hence, they best expose it to severe tests prior to publication. But each of us is also, to some extent, a prisoner of our own conceptual system. We take some things so much for granted that it never even occurs to us to question them. We also have our own career interests. On occasion scientists have refuted the very views for which they are famous, but not often. More often the really severe testing comes from one's opponents. It is also the case that one's opponents are liable to have different though equally unnoticed presuppositions. The self correction so important in science does not depend on scientists being totally unbiased or having no career interests, but on other scientists having different perspectives, not to mention career interests. Scientists working outside your own research group are hurt if they adopt any of your mistaken views, but more importantly, they are also in a better position than you are to expose them to severe tests. *Their* career interests are not damaged if *your* views are refuted.

The preceding has concerned the effects of intra- versus intergroup lying. The effects of intra- versus intergroup stealing are just as dramatic. One's closest research associates are most vulnerable not only to the effects of the quality of one's research but also to have their findings appropriated. As science is now structured, scientists need not make their work available to other scientists until they publish and, hence, establish intellectual ownership. The major exception is the refereeing process. It is possible for a referee to read a manuscript, extract an idea, and get it published in time to gain priority. Because manuscripts submitted for publication are well down the road to completion, stealing of this sort is relatively difficult. Research proposals, to the contrary, are supposed to be largely prospective in nature. They are supposed to describe current research and likely avenues for future work. Hence, reviewers of research proposals have access to the plans of their closest competitors early enough so that they might well be able to gain credit for them before their authors can. As a result, scientists have devised various techniques to obtain research money without giving their competitors too much of an edge. But there is no way to hide one's work from others in one's own research group. Priority disputes between individuals working in relative isolation from each other are vitriolic enough. Those among scientists belonging to the same research group are even more devastating.

Lakatos (1971) has suggested that one way to decide between different views of science is by how many characteristics of science flow naturally from it and how many remain anomalies. The key unit in science as far as Lakatos is concerned is the research program. Research programs in turn are evaluated as to whether they are progressing, stagnating, or degenerating. Because these features of research programs are a function of recognized accomplishments, priority disputes between advocates of different research programs are a matter of "rational interest," while those between scientists promoting the same program result merely from "vanity and greed for fame." Although no one has collected systematic data on this topic, my impression is that priority disputes are as common among scientists working in the same as in different research programs. All of the former are, for Lakatos, anomalies. On my view of science, two processes are involved: individual inclusive fitness as well as intra- versus interdemic selection. Scientists should feel a strong allegiance to their allies for a variety of reasons, some of them nobler than others. Among the less noble is that one's own self-interest is tied up with the interests of one's co-workers. As long as the effects of individual and demic selection coincide, no problems need arise. However, when an individual perceives that his or her individual inclusive conceptual fitness is decreasing in part because of his or her participation in a particular scientific deme, internal friction is guaranteed to increase. On my view, priority disputes both within and between research programs are equally a matter of rational interest.

One might grant that the preceding is a reasonably accurate description of how science as we know it happens to function but nevertheless lodge two objections: first, that this structure is just an historical accident stemming from the general characteristics of the societies in which science happened first to emerge, and second, that science would fulfill its traditional goals even better if its structure were changed. Science did arise in the West in highly competitive, individualistic societies in which property rights were paramount and then diffused to other countries, some of which were, at the time, organized quite differently. If only science had been able to emerge on its own in these other countries, it might have exhibited very different properties than those that happen to characterize it now.

This hypothesis is certainly plausible. I have only two responses to make to it. First, science arose independently several times in the West. In several instances, most notably the French Academy, considerable effort was expended to organize science along more genuinely cooperative lines so that scientists could work in relative anonymity to promote the general good. In every case these efforts failed and were replaced by the system I have described (Hull 1985c). Appeals to work diligently for the general good have never proven to be powerful enough. The greatest strength of science as it is now organized is that it harnesses our "baser" motivations for more "lofty" goals. As the members of the French Academy officially acknowledged in their revised constitution in 1699, in the future instead of working in consort, each member of the Academy "shall endeavor to enrich the academy by his discoveries and improve himself at the same time."

If biology has anything to teach us about functional systems it is that there are always many ways to skin a cat. Functional equivalents are pervasive. Perhaps dozens of different ways exist for the organization of science which would promote our understanding of the natural world. So far only the barest sketches have been presented for these alternative forms of organization, and the few pilot studies that have been conducted have failed decisively. Perhaps the way that science is now organized is far from ideal. It may even be the case that in certain areas competition has become so fierce and the numbers of people involved so large that traditional mechanisms are breaking down. Biomedical research is the best example of science at its worst. Whether these rogue areas of science can be brought back into the fold by more careful attention to traditional mechanisms or whether an entirely new system must be introduced, I for one cannot even begin to guess (for one recent example, see Connor 1987).

SELECTION PROCESSES

Thus far, I have argued the science is a function of the interplay between cooperation and competition for credit among scientists. I have yet to say anything about its selective character. The literature on selection processes in biology has been plagued by a systematic ambiguity in the phrase "units of selection." Some insist that genes are the primary focus of selection because they are the entities that pass on their structure largely intact from generation to generation. Others insist that organisms are the primary units of selection because they are the entities that interact with the environment in such a way that genes are replicated differentially. Still others view the interplay between these two processes as "selection" and define "fitness" accordingly (Brandon and Burian 1984).

For the purposes of a general analysis of selection processes, terms such as "gene," "organism," and "species" are not good enough. Instead more general terms are needed. More general terms are needed if conceptual change is to be viewed as a selection process, but they are also needed in biological contexts as well. If the traditional organizational hierarchy of genes, cells, organs, organisms, colonies, demes, populations, and species is taken as basic, then in point of fact the focus of "selection" wanders from level to level. More than this, traditional entities do not function in the immune system the way that they do in biological evolution, and it is as much a selection process as is biological evolution. For example, selection as it functions in immune reactions takes place entirely in the space of a single generation, and the results of particular selection regimens are not passed on genetically. Selection in immune reactions is ontogenetic, not phylogenetic.

Because such traditional entities as genes, organisms, and species do not consistently fulfill the same roles in biological evolution, not to mention immune reactions and conceptual change, more general units are needed, units that are defined in terms that are sufficiently general to apply to all sorts of selection processes. My suggestions for these units and their definitions are as follows:

replicator — an entity that passes on its structure largely intact in successive replications

interactor — an entity that interacts as a cohesive whole with its environment in such a way that this interaction *causes* replication to be differential.

With the aid of these two technical terms, selection can be characterized succinctly as follows:

selection — a process in which the differential extinction and proliferation of interactors *cause* the differential perpetuation of the replicators that produced them.

Replicators and interactors are the entities that function *in* selection processes. Some general term is also needed for the entities that are produced *as a result of* replication at least and possibly interaction as well:

lineage — an entity that persists indefinitely through time either in the same or an altered state as a result of replication.

In order to function as a replicator, an entity must have structure and be able to pass on this structure in a sequence of replications. If all a gene did was to serve as a template for producing copy after copy of itself without these copies in turn producing additional copies, it could not function as a replicator. Although genes are well adapted to function as replicators, it does not follow from the preceding definition that genes are the only replicators. For example, organisms also exhibit structure. One problem is the sense in which organisms can be said to pass on their structure largely intact. For example, changes in the pellicle of a paramecium are passed on directly when the organism undergoes fission. From the human perspective, populations do not seem to exhibit much in the way of structure, but population biologists recognize something that they term "population structure." If populations can pass on this structure during successive replications of these populations, then they too might function as replictors (Williams 1985).

Many cohesive wholes exist in nature, but only a few of them function in selection processes. Hence, only a very few count as interactors. In order to function as an interactor, an entity must interact with its environment in such a way that some replication sequence or other is differential. Organisms are paradigm interactors. They are cohesive wholes, they interact with their environments as cohesive wholes, and the results of these interactions influence replication sequences in such a way that certain structures become more common, while others become rarer. However, many other entities also function as interactors. For example, genes not only code for phenotypic traits, they themselves have "phenotypes." DNA is a double helix which can unwind and replicate itself. In doing so it interacts with its cellular environment.

In the beginning, one and the same entities had to perform both functions necessary for selection. Because replication and interaction are fundamentally different processes, the properties which facilitate these processes tend also to be different. None too surprisingly, these distinct functions eventually were apportioned to different entities. When a single structure or entity must perform more than one function, it usually performs none of them very well. Too many compromises have to be made. Interaction occurs at all levels of the organizational hierarchy, from genes and cells, through organs and organisms, up to an possibly including populations and species.

As I have characterized it, selection is an interplay between two processes — replication and interaction. Both processes taken separately and the interplay between them are causal processes. As a result, drift does not count as a form of selection — as it should not. An entity counts as an interactor only if it is functioning as one in the process in question.

Thus, if changes in replicator frequencies are not being caused by the interactions between the relevant interactors and their environments but are merely the effects of "chance," then the changes are not the result of selection. Drift is differential replication in the absence of interaction.

Many entities persist indefinitely through time. Of these, some change while some do not. However, the only entities that can count as lineages in the technical sense which I am proposing are formed by sequences of replicators. Hence, on my usage, "lineage" is inherently a genealogical concept. For example, the solar system has changed through time. Nevertheless, it does not count as a lineage because no replication was involved in such changes. The general notion is that of an historical entity — a space-time worm. Lineages are historical entities formed by replication. Differential perpetuation caused by interaction is not necessary for something to count as a lineage. In fact, differential perpetuation itself, regardless of its causes, is not even necessary for something to count as a lineage. However, when the interplay between replication and interaction causes lineages to change through time, the end result is evolution through selection.

Both genes and organisms form lineages. In most cases gene lineages are wholly contained within organism lineages. According to more gradualistic versions of evolutionary theory, species do not *form* lineages. Instead they themselves *are* lineages. But nothing about the evolutionary process requires that it be gradual. It might be the case that species are incapable of indefinite change and that speciation is always saltative. If so, then particular species themselves do not count as lineages. Instead, successions of species form lineages. It might be the case that not all organisms belong to species. For example, if significant gene exchange is necessary for species to exist, then they did not exist for the first half of life on Earth and are still absent among a significant proportion of organisms living today. Even though species are not a necessary consequence of replication, lineages are.

THE ROLE OF INDIVIDUALITY IN SELECTION

By now, one thing should be clear. Everything involved in selection processes and everything that results from selection are spatiotemporal particulars — individuals. Both replicators and interactors are unproblematic individuals. To perform the functions they do, they must have finite durations. They must come into existence and pass away. Replicators must exhibit structure, and interactors must interact with their environments as cohesive wholes. These are the traditional characteristics of individuals. That individuality is at the heart of the selection process can be seen by the frequency with which biologists who want to argue that species can be selected begin by arguing that they have the properties usually attributed to individuals, ordinary perceptions and conceptions notwithstanding (Eldredge and Gould 1972). Conversely, those who argue against species selection begin by arguing that species lack these very characteristics. In fact, those who argue that not even organisms can function as units of selection begin by casting doubt on their status as individuals, superficial appearances notwithstanding (Dawkins 1976).

Lineages are also individuals but of a special sort. In order to function as a replicator, an entity can undergo minimal change before ceasing to exist. In order to function as an interactor, an entity can undergo considerable but not indefinite change. Lineages are peculiar in that the organization which they exhibit is sufficiently loose so that they can change indefinitely through time but sufficiently tight that the effects of selection are not lost. Thus, any entity that can function either as a replicator or as an interactor cannot function as a lineage because they are too tightly organized. Conversely, any entity that can function as a lineage cannot function as a replicator or as an interactor because it lacks the requisite internal cohesiveness. All sorts of factors can enhance or destroy the cohesiveness of a lineage, but genealogical connectedness through time is essential.

All of the preceding may seem excessively fluid, but this fluidity is dictated by the nature of living entities. According to what we know of the functioning of the genetic material in both autocatalysis and heterocatalysis, genes are anything but beads on a string. Organisms come in a dismaying variety of forms, some well-integrated, others not. Even though colonial forms of organization depart significantly from our vertebrate perspective, they are widespread. Some organisms form colonies; others not. Some species are composed of demes; others more homogeneous, and so on. Any adequate theory of biological evolution must apply to all organisms, not just to well-integrated sexual organisms. As strange as the notions of tillers and tussocks, genets and ramets may sound to a zoologist, plants evolve too (Jackson, Buss, and Cook 1986).

What is more, entities do not stay put as they change though time. For example, in certain colonial wasps, each hive initially includes several queens. As time goes by all are killed off but one. As a result, the focus of selection expands from the individual organism to the entire hive. In a large, genetically heterogeneous species in which crossover is common, very small segments of the genetic material are likely to be the primary replicators, but when such a species goes through a populational bottleneck, entire genomes may come to function as single replicators. Similar observations hold for lineages. In any system that is evolving though selection, a point occurs at which the genealogical fabric is rent and networks become trees. At whatever level in the traditional hierarchy that this occurs, the resulting entities are lineages. For example, among strictly asexual organisms, no lineages exist that are more inclusive than organism lineages. Among sexual organisms, some form lineages no more inclusive than sequences of populations. In some species gene exchange may be sufficiently extensive and sustained to integrate entire species into lineages. Among some plants, gene exchange among traditional taxonomic species may be sufficient to integrate them into a single lineage (Mishler and Donoghue 1982).

All of the preceding may seem not only too fluid but also wrong-headed. Genes and organisms are unproblematic individuals, while species are just as unproblematic classes. After all, the terms "organism" and "individual" are interchangeable. First off, common usage may be the place where all investigations are forced to begin, but it cannot be where they end as well. At times, ordinary usage is misleading and must be modified. Secondly, all organisms may count as individuals according to ordinary usage, but not all individuals are organisms. "Individual" refers to a much broader spectrum of entities than organisms even in ordinary usage, including stars, continents, buildings, and nations. Furthermore, when one tears one's attention away from vertebrates and looks at all living creatures, many organisms turn out not to be paradigmatic individuals. Many biologists do assume uncritically that all organisms are individuals and vice versa, but they are wrong to do so.

Such issues to one side, some of my readers may take the distinction between spatiotemporally localized individuals and spatiotemporally unrestricted classes (or sets) to be of no consequence. Once again, terminology intrudes. I am not in the least interested in the terms that are used to mark this distinction. I take the distinction to be important, not the terms. The distinction is important because it has characterized science throughout its existence. Whether or not scientists have been mistaken to do so, one of the fundamental goals of the most important scientists in the history of science has been the discovery of natural regularities that apply to any entities whatsoever, just so long as they meet certain conditions. According to some, the claim that Moses wandered in the Sinai or that all the coins in my pocket are dimes differs in no important respects from Newton's law of universal gravitation. If so, then we have all been behaving in extremely inappropriate ways because we pay very little attention to those who come up with trivial claims of the first sort and heap fame and fortune on those few scientists who produce statements of the second sort. As difficult as it may be to present a totally satisfactory analysis of the distinctions between singular statements, accidentally true universal generalizations, and laws of nature, these distinctions are fundamental to our understanding of science.

Implicit in the preceding discussion is the distinction between two sorts of entities: those that are spatiotemporally restricted in the relevant sense and those that are not. Perhaps for philosophical purposes, Moses, all the coins in my pocket, *Dodo ineptus*, gold, and bodies with mass are all equally "sets." If so, then I am forced to distinguish between two sorts of sets: those that must be spatiotemporally restricted and localized to perform the roles they do in the natural processes in which they function, and those that must be spatiotemporally unrestricted to perform their quite different roles. To function as a replicator or an interactor, an entity must be spatiotemporally localized and cohesive, while lineages must be less tightly-organized but just as spatiotemporally localized.

Once the preceding distinctions are made, decisions as to which entities belong in which category are largely a contingent matter. Do all genes and only genes function as replicators in the evolutionary process? The answer is clearly no. Do all organisms and only organisms function as interactors in the evolutionary process? Once again, the answer is clearly no. How about species? Given the groupings of organisms commonly considered species by systematists, do all and only species form lineages as a result of biological evolution? Although others might think otherwise, I think that the answer to this question is just as clearly no. Traditionally, particular species such as *Dodo ineptus* have been treated not simply as classes but as natural kinds akin to gold and triangle. However, if species are considered as chunks of the genealogical nexus, then they are as much spatiotemporally restricted and localized as are organisms and genes. Perhaps one might not want to term species "individuals," perhaps one might want to introduce a third category to the traditional distinction between individuals and classes for lineages (assuming species count as lineages), but species cannot be construed as spatiotemporally unrestricted regardless of whether they function in biological evolution as replicators or interactors, or only result from the action of selection processes operating at lower levels (Ghiselin 1974a, Hull 1976, Mayr 1987). Common sense to one side, genes and organisms are not unproblematic examples of spatiotemporally localized individuals, and species are not unproblematic examples of spatiotemporally unrestricted classes.[2]

SCIENCE AS A SELECTION PROCESS

All of the preceding concerns biological evolution. If any regularities are to be found in evolutionary processes, general conceptions such as those I have set out are necessary. Perhaps my particular conceptions will turn out not to be good enough, but if both biological evolution and the reaction of the immune system to antigens are to count as selection, then traditional conceptions will not do. More general conceptions are necessary even in strictly biological contexts. My concepts have the added virtue that they are sufficiently general to apply to conceptual evolution as well, in particular to conceptual selection in science. None too surprisingly, the replicators in science are elements of the substantive content of science – beliefs about the goals of science, the proper ways to go about realizing these goals, problems and their possible solutions, modes of representation, accumulated data reports, and so. Scientists in conversations, publications, and lectures broach all of these topics. These are the entities that get passed on in replication sequences in science. Included among the chief vehicles of transmission in conceptual replication are books, journals, computers, and of course human brains. As in biological evolution, each replication counts as a generation with respect to selection.

As they function in the production of proteins, genes are organized into reasonably discrete, hierarchically organized functional units. Because crossover does not respect the boundaries of these functional units, the chunks of the genetic material that are transmitted intact are highly variable. In biological evolution, these variable chunks of the genetic material are the primary replicators. According to some authorities, organisms and even possibly gene pools can also function as replicators in the evolutionary process. In any case, genes both have phenotypes of their own and code for more inclusive phenotypes which influence their perpetuation by means of their relative success in coping with their respective environments. Conceptual replicators interact with that portion of the natural world to which they ostensibly refer no more directly than do genes with their more inclusive environments. Instead they interact only indirectly by means of scientists. Scientists are the ones who notice problems, think up possible solutions, and attempt to test them. They are the primary interactors in scientific change.

Conceptual replication is a matter of information being transmitted largely intact from physical vehicle to physical vehicle.[3] In addition to accuracy of transmission, these vehicles have two important characteristics – their duration and how active they are. Certain vehicles are much more ephemeral than others. The spoken word is extremely transitory, human beliefs incorporated into individual brains can last for longer periods of time, but people die. If a belief is to survive, it must be replicated. In books and journals, ideas find a much more durable medium. They can enter into a replication series and then lie fallow for generations, until someone else happens to stumble upon them to initiate a new series. Or they can be passed on unnoticed, the way that certain atavistic genes are, until they begin to function again. Scientists' brains can serve as vehicles for replication sequences, but scientists themselves are anything but passive vehicles for such sequences. They also function as interactors. Without scientists, no conceptual replicator could ever be tested, and testing is essential to science. As a shorthand expression, we frequently talk about the meanings of words and sentences, and in developed languages there is some point to such circumlocutions, but ultimately people are the entities who mean things by what we say. Individual scientists are the agents in scientific change.

In sum, conceptual replication is a matter of ideas giving rise to ideas via physical vehicles, some of which also function as interactors. Replicators are generated, recombined, and tested by scientists interacting with the relevant portion of the natural world. Because I see a ball accelerate as it rolls down an inclined plane, I come to hold beliefs about the motion of balls as they roll down inclined planes. Something in the non-conceptual world initiated a replication sequence in the conceptual world. These sequences of events in the non-conceptual world are the sorts of causal connections that natural science is designed to discover. Social scientists study the perceptual connections between individual organisms and the rest of the world, including other organisms.

Causal connections exist between scientists and the non-social natural world, but they also exist among scientists themselves. Science is not only a conversation between individual scientists and the natural world but also a conversation among scientists. Causal connections also exist between scientists and their societies at large. For example, throughout its history, science has been carried out in sexist societies. Perhaps this sexism has "infected" science itself. These larger social influences are not always of the sort that need to be "overcome," but when they are, the social organization of science permits it. Because the career interests of individual scientists frequently conflict with each other while coinciding often enough with the manifest goals of science, scientists not only come to notice the effects of broader social interests on their belief systems but also sometimes overcome these effects. Perhaps scientists raised in a sexist society should not be able to notice the sexism latent in their society let alone overcome it, but they do.

Previously I noted that replicators must not only exhibit structure but also pass it on to subsequent replicators. But structure alone is not enough. This structure must count as "information." In the case of selection, biological evolution has been taken to be the literal usage, and the transfer to conceptual change is usually considered, if not dismissed as being "merely" analogical. In the case of information, conceptual transmission is the literal usage, and the genetics context is analogical. Many regularities exist in nature. Planets travel in ellipses, gases expand when heated, molecules are transported differentially through semipermeable membranes, and crystals form very regular shapes. None of these regularities count as "information" in the appropriate sense. The order of bases in molecules of DNA does because of the character of this order, its origins, and its effects.

Much of the structure of DNA is strictly lawful. Given certain general constraints, it has to be the way it is. For example, if the external "backbones" of a DNA molecule are to be kept equidistant down its length, then the "rungs" between these backbones must be kept of equal length. As its name implies, deoxyribonucleic acid is an acid and, as such, must have the general characteristics of acids. These are the features of

DNA which allowed scientists to unravel its general structure. The order of bases in a particular molecule of DNA is quite different. Adenine can bind only with thymine and quanine can bind only with cytosine, but with only minor exceptions, any of the four bases can precede or follow any of the other bases as well as itself. All orders are equally likely from the perspective of physical law. That is why the order that happens to exist can function as a code. With equally minor exceptions, any letter in a language such as English can precede or follow any other letter.

As problematic as the distinction is between those features of a molecule that follow lawfully from the fundamental character of the physical world and those that are contingent, it is nevertheless important. The distinction between a code (or language) and particular messages expressed in that code is similarly quite important. Both natural languages and messages are built up historically and can be used to infer history. Numerous genetic codes were possible when life evolved on Earth. Which code happened to evolve is primarily a function of the particular circumstances that obtained at its origin. Because all terrestrial organisms use the same code with exceptions that are as slight as they are rare, the assumption is that all life here on Earth had a single origin. Although the genetic code is relatively simple when compared to a natural language such as English, it is still sufficiently complicated that the likelihood that exactly the same code could have evolved twice is extremely small. Although all organisms use the same code, quite obviously they contain different messages. Because later messages are modifications of earlier messages, molecular biologists have been quite successful in reconstructing the past history of extant forms of life.

In a sense, any record of the past can serve as information about the past. For example, unusually high concentrations of iridium in thin layers of the geological strata can be used to infer past impacts with meteors, footprints of dinosaurs in the hardened mud of a river bed imply the existence of these organisms at the time that this layer was formed, and masses of charred wood indicate a forest fire. Organisms can learn about the world in which they live either directly by interacting with it or indirectly by observing some other organism do the interacting. Learning from experience in the first instance may initiate a replication sequence but does not itself count as replication. In their investigations, scientists learn about the structure of the empirical world. They record this knowledge in a language of some sort. This characterization of the natural regularity counts as information, but the natural regularity itself cannot without making the notion of information vacuous. The chief exception is knowledge of the genetic code. In this instance, the regularity that initiates a conceptual replication sequence is itself part of a replication sequence.

Scientists have carried the process of learning to its extreme. Students can run experiments themselves or watch their instructors do so. But most learning in science comes from reading or hearing about the activities of others. Only a small portion of what a scientist believes about the world arises by means of this scientist interacting with the relevant phenomena. Each scientist has only a few decades to contribute to science. Time cannot be wasted checking every knowledge claim before using it. Using without testing makes scientific progress possible, but it also increases the possibility that some of these knowledge claims are likely to be mistaken. However, one should not forget that knowledge by acquaintance, no matter how direct, is also far from infallible.

DISANALOGIES BETWEEN BIOLOGICAL AND CONCEPTUAL EVOLUTION

Numerous differences have been alleged between biological and conceptual change. Most have very little substance and can be treated quite briefly. Others have some point and deserve a fuller discussion than I can give here. For example, one frequently hears that conceptual evolution occurs much more quickly than biological evolution. In point of fact, conceptual evolution occurs at an intermediate rate as far as physical time is concerned. Viruses evolve much more quickly than conceptual systems in even the most active areas of research, while large, multicellular organisms evolve more slowly. However, physical time is relevant only to interaction. As far as replication is concerned, the relevant metric is generation time. With respect to generations, conceptual evolution occurs at the same rate as biological evolution — by definition.

Some authors argue that no general analysis of selection processes equally applicable to biological and conceptual evolution is possible because genes are "particulate" while the units in conceptual replication are highly variable and far from discrete. In point of fact, neither biological nor conceptual replicators are all that "particulate." In both cases, the relative "size" of the entities that function either as replicators or as interactors is highly variable and their boundaries sometimes quite fuzzy. If the entities that function in selection processes must all be of the same size and/or be sharply distinguishable from each other, then selection can no more occur in biological than in conceptual contexts.

Another objection that has been raised is that biological evolution is always biparental, while conceptual evolution is usually multi-parental. Once again, this objection is based on a simple factual error. For a large number of organisms, inheritance is biparental; for most it is not. In conceptual evolution, rational agents sometimes combine ideas from only two sources; sometimes from several. Upon casual inspection, polyploidy seems somewhat more common in conceptual than in biological evolution, but that is all. When one switches from the level of individual entities to populations, no significant differences can be found between the two.

At a particular locus, numerous different alleles can coexist in various frequencies. Numerous different solutions to the same problem or versions of the same idea can coexist in conceptual populations.

Upon first glance, cross-lineage borrowing seems much more common in conceptual than in biological evolution. If lineages are defined in terms of replication sequences, extensive cross-lineage borrowing is ruled out by definition. Two lineages can remain distinct in the face of some cross-lineage borrowing, but once it becomes too extensive, the two lineages merge into one. Regardless of commonsense beliefs, gene exchange does occur between groups that are considered different species, and the amounts of gene exchange needed to neutralize any genetic differences between two largely disjoint lineages turns out to be quite small. In short, in biological evolution, extensive cross-lineage borrowing cannot occur because lineages are generated by this very process. When conceptual and social lineages are distinguished in science, extensive cross-lineage borrowing becomes possible, i.e., scientists belonging to different social-defined groups can and sometimes do use each other's work. In such situations, the groups remain socially distinct while their conceptual correlates merge. However, such cross-lineage borrowing in science does not seem to be as extensive as recurrent references to various "syntheses" would lead one to expect. Rarely do conceptual lineages merge without the scientific communities that produced them merging as well. Mergers of both sorts occur in science. They occur in biology as well, especially among plants. So far no one has produced the data necessary to see in which contexts cross-lineage borrowing is more prevalent.

The most commonly cited disanalogy between biological and conceptual evolution is that biological evolution is Darwinian while conceptual evolution is largely Lamarckian. No organism is able to pass on any of the ordinary phenotypic traits that it acquired during the course of its existence to its progeny, but some organisms can pass on what they have learned about their environment through social learning. As often as such observations are repeated, no one goes on to explain at any length what they mean. No one claims that conceptual evolution in science is *literally* Lamarckian, as if the basic axioms of quantum theory are somehow going to find their way into our genetic make-up. If conceptual entities are taken to be phenotypic traits, then conceptual evolution is not literally Lamarckian because changes in these traits leave genes untouched. Ideas are transmitted but not inherited. If simple transmission is sufficient for Lamarckian inheritance, then a mother giving her baby fleas counts as Lamarckian inheritance. Taken *metaphorically*, conceptual evolution is still not Lamarckian because ideas (or memes) are held to be the analogs of genes, not characters. If anything, conceptual evolution is an instance of the inheritance of acquired memes, not characters. We do learn from experience and pass on this knowledge socially, but I fail to see why these processes should be considered "Lamarckian" in either a literal or a metaphorical sense. On the literal interpretation, ideas count as acquired characters, but the transmission is not genetic. On the metaphorical usage, ideas count as analogs to genes, not characters. Although the genotype-phenotype distinction can be made in the context of conceptual change, the net effect is that the analogs to phenotypes are not inherited. In the absence of anything like the inheritance of acquired characters, I think that characterizing conceptual change as "Lamarckian" leads to nothing but confusion.

As far as I can see, the *only* sense in which conceptual evolution is Lamarckian is in the most caricatured sense of this much-abused term, i.e., it is intentional. Just as giraffes increased the length of their necks by striving to reach leaves at the tops of trees, scientists solve problems by trying to solve them. Science is intentional, in fact it is as intentional as any human activity can get. We learn about the natural world by contriving to interact with it. For some, the gulf that separates intentional acts from the rest of nature is so wide and deep that no comparisons are possible. I do not share this conviction, but I have no in-principle arguments that are liable to touch those who wish to insulate the behavior of intentional agents from the sort of principles that apply to the rest of the natural world. All I can do is to point out some of the consequences of taking this distinction as primary and all others as secondary. For example, in the *Origin of Species*, Darwin reasoned from the known effects of artificial selection to the possible effects of natural selection. But artificial selection is intentional. Perhaps plant and animal breeders cannot produce mutations at will, but they do consciously choose those organisms to breed that exhibit traits which they find desirable. If reasoning from artificial selection to natural selection is totally illicit, then Darwin's main argument in the *Origin* is one gigantic blunder. Similarly, any extension from controlled experiments in science to the rest of nature is illicit.

I also do not think that the role of intentionality in scientific contexts is actually at the root of what bothers critics of any attempts to provide a single analysis of "selection" that applies equally to biological and conceptual evolution. Scientists do strive to solve problems. They both generate novel ideas and select among them. Right now genetic mutations occur by "chance." However, in the very near future, biologists will be able to generate any genetic mutations they see fit. When that occurs, intentionality will play the same role in both biological and conceptual change. I doubt that in such an event critics will instantly become converted. If I am correct in my guess, then the role of intentionality in generating novelty must not have been all that important of an objection in the first place.

For my part, I think a better way of classifying selection processes is between, first, those that are gene-based and those that are meme-based and only then worry about the complications which ensue from intentionality. Some gene-based selection is intentional; most is not. Some meme-based selection is intentional; most is not. On this classification, artificial and natural selection are fundamentally the same sort of phenomenon. On this same classification, the sort of rational selection of beliefs in which people engage (when all else fails) and all the semiconscious and unconscious selective retention that characterizes how human beings acquire their beliefs are also fundamentally the same sort of phenomenon. Once one looks at science as a whole, paying attention not just to the rare scientist who happens to make a major contribution to science but to the vast army of scientists who have no discernible impact on science, the effects of intentionality do not look so massive. If scientists did not strive to solve problems, the frequency with which they succeed would no doubt decrease, but it is already so low that the differences would be difficult to discern. All scientists are constantly striving to solve problems. Few do. Of those who do, only a very few are noticed. There may well be a difference in kind between intentional and non-intentional behavior, but it is not a difference in kind that results in much of a difference in degree.

One difference between biological and conceptual evolution is that in biology genes make genes. In the most primitive circumstances, genes were also probably the only interactors. Eventually, however, they began to produce more inclusive entities that could promote replication by interacting with their more inclusive environments. Conceptual replicators do not, on their own, produce copies of themselves. They do so only via their most important interactors — individual scientists. Thus, in scientific change, scientists are the chief agents in both replication and interaction. However, on my analysis, this difference is not sufficient to preclude a single analysis applying equally to both.

Another apparent difference between biological evolution and conceptual change is that biological evolution is not clearly progressive while in certain areas, conceptual change gives every appearance of being progressive. At a glance, biological evolution appears to be as clearly progressive as conceptual evolution in the most advanced areas of science, but appearances are deceptive. Thus far biologists have found it surprisingly difficult both to document any sort of biological progress in the fossil record and to explain what it is about the evolutionary process that might lead phylogenetic change to be progressive (for the most convincing data to date, see Signor 1985).

Conceptual development in certain areas of human endeavor, especially in certain areas of science, gives even a stronger appearance of being progressive. Although science is not progressive in the straightforward way that earlier enthusiasts have claimed, sometimes later theories are better than earlier theories even on the criteria used by advocates of the earlier theories. Science at least appears to be more clearly progressive than biological evolution. Of greater importance, we have good reason to expect certain sorts of conceptual change to be progressive.

Intentionality is close to necessary but far from sufficient in making conceptual change in science progressive. 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. The frequency of success in science is quite low. Even so, one should expect that, on average, scientists should solve the problems which they are trying to solve more frequently than those which are only at the periphery of their attention. At the very least, the intentional character of science should speed it up. However, intentionality is far from sufficient in explaining the progressive character of science. 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 behavior can have a direction in a global sense only when the goal stays put.

Whenever the conditions are right, evolution by means of natural selection occurs. The global goal of natural selection may well be increased adaptation, but for particular lineages, the contingencies to which successive generations of organisms must adapt keep changing, not because genetic variation is "blind," not because natural selection is non-intentional, but because so many of the aspects of the environment to which organisms must adapt keep changing. 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.

CONCEPTUAL INTERACTION

In biological evolution, replicators pass on their structure largely intact. Some of this structure counts as "information." Through interaction this information is translated into phenotypes at a variety of levels. One reason why selection processes are so complicated is that they evolve via an interplay between two subsidiary processes (replication and interaction) occurring at a variety of levels. Because so little of the information encoded in the relevant replicators is ever realized, selection processes are extremely particularistic and idiosyncratic in their effects. For instance, genes do not literally code for traits. Given any genome, a wide spectrum of traits could eventuate depending on the sequences of environments confronted. Of all the phenomes a particular genome could have produced, only one is produced. A genome that might well have proven to be extremely fit in a wide variety of environments is extinguished because it does not happen to find itself in one of these environments. Even when a single genotype is expressed clonally in numerous genomes, only a tiny fraction of the phenomes that could have been produced actually are produced. Hence, genotypes never get to show "all their stuff." In short, translation entails a tremendous loss of information.

Similar observtions hold for conceptual evolution. Scientists do spend an appreciable amount of time testing their views, but as everyone now acknowledges, scientific theories are always grossly underdetermined by anything that might be considered data. Given any set of observations, the number of alternative theoretical explanations that might be generated to account for them is limited only by the ingenuity and good sense of scientists. Of all the possible explanations that could be offered, only a small fraction actually are ever offered. Conversely, of all the observational implications of any one version of a particular theory, only a small percentage actually are ever made, let alone tested. As a result, chance plays a large role in which versions of which theories ever become prominent. A particular theory might generate considerable attention even though it has serious defects because the first observations made happen to be among the relatively few that actually conform to it. Conversely, a theory that in retrospect has much to recommend it might be rejected because the first observations made in testing it happen to be among the relatively few that are at variance with it. The slippage which exists between theories and data is only exaggerated by the number and variety of compromises that must be introduced in order to "operationalize" a theory so that it can be tested.

Biologists were not prepared for the amount of genetic variation that characterizes biological species. More often than not, greater heterogeneity exists within a species than between it and its closest evolutionary congeners. The same heterogeneity characterizes science. Praise for conceptual pluralism in science is currently fashionable. Conceptual pluralism is necessary if science functions as a selection process, but so is conceptual pruning. Taken literally, the maxim that anything goes would be lethal to science. Scientists are constantly generating different combinations of the conceptual tools which have been willed to them by previous generations. At times, in fact quite rarely, genuine novelties are also introduced. At a distance, much of this conceptual heterogeneity is obscured by the deceptive appearance of terminological conformity. Scientists are as terminologically conservative as they are semantically flexible. Under cover of the same terms, scientists working in the same research program frequently hold very different views.

Kuhn (1970) attempted to make his notion of a "paradigm" more operational by reference to scientific communities. Paradigms are those things that members of the same socially-defined scientific community share. The trouble is that if one actually defines a group of scientists in terms of their professionally relevant relations, the resulting groups tend to be conceptually quite heterogeneous. Those of us who study science tend to find this heterogeneity quite disconcerting. We are convinced that the only way that people can cooperate is for them to agree with each other at least over fundamentals. Every scientific theory *must* have an essence, and every scientist working in the same research program *must* accept these essential tenets.

I do not know about people in general, but scientists seem to be able to cooperate with each other even when they are in fundamental disagreement. They do so, in part, by playing down these differences. When asked, scientists insist that the conceptual system that they are developing can be characterized by a set of fundamental propositions about which there is universal agreement. However, when the members of a socially-defined research group list the fundamental principles of their research program, they present different lists. Even when some items on this list are terminologically the same, the intent frequently varies. For example, a group of systematists might all agree that all higher taxa must be monophyletic but mean very different things by "monophyletic."

Much of this heterogeneity is lost as particular research programs "harden." In retrospect a particular program may well appear to have had an essence or Lakatosian "hard core," but as Lakatos himself noted, hard cores can be recognized only in retrospect. While the selection process is going on, conceptual systems are heterogeneous — as they must be if science is a selection process. When scientists say that the research program they are working on has an "essence," they mean the views that they individually happen to hold at the moment. *The* essence of a Darwinian view of evolution is *their* view, the benighted opinions of others who consider themselves Darwinians notwithstanding. Sooner or later such dogmatism does triumph. One version will come to be accepted as *the* version, and all the variation which characterized the program in its active period will be ignored or discounted. The challenge then is to provide a method for making sense of scientific change in the face of all this conceptual heterogeneity.

THE TYPE SPECIMEN METHOD OF REFERENCE

On the account of scientific change that I have sketched above, groups of scientists must be distinguished from the conceptual systems that they produce and the two followed separately. Both form internally heterogeneous lineages that can change through time. A research group can persist while old members leave and new members join. Conceptual lineages are no less heterogeneous. At any one time, they can contain contradictory elements, and a particular statement-token can give rise through successive replications to a statement-token that contradicts it. In the face of all this heterogeneity and change, how can those of us who study science make any sense of scientific change? How can we individuate and refer unambiguously to a particular lineage, either social or conceptual? One possible solution to these problems is the appropriation of a method devised by systematists through the centuries to handle parallel problems with respect to biological species.

When systematists come across what they take to be a previously unknown species, they pick a specimen, any specimen, and designate it as the type specimen. In spite of the connotations of this term, the type specimen need not turn out to be in any sense "typical." All a type specimen does is determine to which species a name applies. No matter how aberrant a type specimen may turn out to be, it belongs to a particular species and to no other. It is one node in the genealogical nexus. No matter how the boundaries of a particular species are reworked, the species that includes the type specimen must be called by the name that the type specimen bears. Both the name and the type-specimen are passed down through the generations from systematist to systematist.

As similar as the type-specimen method in systematics may look to the theory of rigid designation (Kripke 1972, Putnam 1973), it differs from it in several important respects. In systematics fictitious baptisms play no role, and systematists treat link-to-link transmission chains seriously. Systematists actually perform the sort of historical inquiry necessary to trace names back to their sources and make their decisions accordingly. Priority is one of the key elements in their codes of nomenclature. But most importantly, biological species are chunks of the genealogical nexus and as such count as spatiotemporal particulars and not classes or kinds. As a result, their names are best construed as being proper, not general. Because proper names have been treated traditionally as "rigid," one need introduce no new theory of reference to accommodate them. Like an organism such as Moses, a species such as *Dodo ineptus* has a beginning, middle, and end. A name can be attached to it rigidly during any time-slice of its existence. If one chooses, one can apply the same name to a lineage throughout its existence, no matter how much it might change, or subdivide it, giving each subdivision a separate name. For example, one might choose to subdivide an organism into sequential stages and name each, just as one can subdivide a gradually evolving lineage into successive chronospecies, giving each a different name, but the logic of the situation remains the same. The lineage is basic; the characters describing the entities that are part of the lineage are secondary.

The type-specimen method works so well for historical entities because *both* the entity being named *and* the subsequent link-to-link transmission of its name form historical entities that can be traced independently of meaning change to see if, in the past, they intersect in the way claimed. Did someone named "Moses" exist at the time and place claimed in the Bible? Did the dodo ever exist on Madagascar and only recently go extinct? These questions have answers which are independent of a whole variety of other considerations. "Moses" would still designate Moses even if he did not do many of the things attributed to him in the Bible, and "dodo" would still designate the dodo even if we happily discovered a population of these birds still alive in some remote valley somewhere.

The type-specimen does not work quite so well for terms that are genuinely general because the entities referred to are not themselves historical entities. The substances gold and water can exist anywhere in the universe whenever the conditions are right. They can be named numerous times both here on Earth and elsewhere. Because languages here on Earth have histories, and these histories are to some extent interconnected, sometimes link-to-link transmissions of term-tokens can be traced back through time, and sometimes they converge, but they need not. "Wasser" and "water" may well have the same terminological ancestor, but the ultimate reason why the plethora of terms of denoting this substance are held to denote the same substance is similarity in structure of this substance.

However, if conceptual change in science is taken seriously as a selection process, then something like the type-specimen method can handle it, but this revised theory depends crucially on organizing term-tokens into term-trees solely on the basis of transmission. In these trees, the structure of the tokens can change, e.g., tokens of "pangen" can be transcribed as "pangene" and then "gene." Even so, these term-tokens belong to the same token-tree. Even the character of the events that initiate the use of term-tokens can change. For example, De Vries used tokens of the term "pangen" only in certain circumstances, circumstances significantly different from those that elicited the tokens of the term "gene" from Johannsen. Present-day scientists use this same term under an even greater variety of situations. What binds all these term-tokens together is that they all belong to the same term-token tree. Actual baptisms did occur, and subsequent uses actually form intersecting trees (for further discussion of the application of the type-specimen method to conceptual change in science, see Hull 1983b; for objections, see Mayr 1983).

One can group organisms together in a variety of ways. Which ways are preferable depend on the use that can be made of these groupings. For example, one can group organisms into those that reproduce sexually and those that do not. If these groups are treated as being genuinely general, then the terms referring to them might well function in laws of nature. Or one might group organisms according to descent. If so, then these groups must be spatiotemporally restricted, and the terms that denote them cannot function in spatiotemporally unrestricted generalizations. However, at the very least, grouping by descent serve to recognize the entities that result from natural selection. Some of these groups themselves might even function in selection processes. The distinction is between "Carnivora" and "carnivorous." The former refers to a chunk of the genealogical nexus; the latter to a group of organisms that share the ability to eat and digest meat. The two groups are far from coextensive. Not all species that belong to Carnivora are carnivorous, and many species of carnivorous organisms do not belong to Carnivora. Both sorts of groupings have their function in biology, but these are different functions.

But, one might object, why not equivocate between these two different uses of the term? Certainly that is what most people do most of the time in ordinary language. Sometimes "Baroque" is used as if it applied just to a particular period in human history, sometimes in a more general sense. Why not term a chair or building constructed today "Baroque"? Insisting that this term applies only to a particular time and place is being too "monistic." Although distinguishing between Tiffany lamps and Tiffany-type lamps might seem overly pedantic, it is no more pedantic than distinguishing between characters such as eyes (in the sense of any organ that can be used to discern light) and eyes as evolutionary homologies. Vertebrate eyes and cephalopod eyes are not the "same" character. In the context of selection processes, these two senses of "same" must be distinguished. Some genes are identical, some identical by descent, and the roles in population genetics of these two sorts of identity differ (Hull 1986).

Parallel distinctions exist for term-tokens. They can be grouped into types solely on the basis of something such as "similar meanings" regardless of genesis. If so, then any token of this term-type that has the appropriate meaning belongs to that type. These term-tokens can appear anywhere at any time. Although it is unlikely that the same term-type should be coined independently to refer to the same thing, it would make no difference if such an unlikely occurrence did take place. In fact, if each term-token of a term-type were generated *de novo* with each utterance, it would make no difference. The notion of "sameness of meaning" has proven to be extremely slippery. Even so, if sameness of meaning is what groups term-tokens into term-types, then genesis is irrelevant.

If one wants to treat conceptual change as a selection process, then term-tokens must be grouped into lineages and trees by means of transmission. Link-to-link transmission must be taken literally. In these replication sequences, term-tokens themselves can change (e.g., "pangen" can become "pangene"). They can also change the means by which they are connected to their referents (e.g., additional operational "definitions" for "gene" can be introduced). As strange as this way of grouping term-tokens may seem, it is necessary if conceptual systems are to evolve by means of selection.

Both ways of grouping term-tokens have their function in science. Within the context of a particular controversy in science, the causal connection of term-tokens in replication sequences is crucial. Term-tokens are the things that are being differentially perpetuated. Anyone who wants to understand scientific change at the local level must order term-tokens into trees. As "irrational" as it may seem, scientists evaluate claims in terms of their genesis because of the influence of both conceptual inclusive fitness and the demic structure of science. Two instances of the same statement-type are evaluated differently if they happen to be part of two different lineages. One might be rejected, resulting in that conceptual lineage going extinct. The other might be accepted and proliferate until it is universally accepted. However, those involved in these selection processes intend their usage to be general. They intend to transmit term-types, but all they actually transmit are term-tokens which are immediately interpreted as types.

Thus, conceptual change in science viewed as a selection process incorporates a systematic ontological equivocation. Term-tokens are tested and transmitted locally but interpreted globally as types. Term-tokens are simultaneously part of spatiotemporally restricted term-trees and instances of spatiotemporally unrestricted term-types. Each generation of scientists intends for their conceptual systems to be generally applicable and universally accepted, but in each generation only a very small percentage of instances of these systems gets passed on, and the version of a particular conceptual system that eventually comes to prevail may well not be the one that early scientists intended. As strange as the distinction between identity and identity by descent may appear, it permeates all evolutionary biology. For many processes, identity by descent is required; for others it is not. Similar genes or traits behave similarly in similar situations, differences in genesis notwithstanding. If bottlenecks are as important in biological evolution as many evolutionary biologists claim they are, then identity by descent is crucial in the evolutionary process. Most change occurs when a population goes through such a bottleneck. Similarly, if small research groups are as important as some students of science claim that they are, then similar observations hold for the role of identity by descent in conceptual change in science.

CONCLUSION

Every explanation takes certain things for granted and explains other things in terms of them. I have taken for granted that scientists are by and large curious about the world in which they live and desire credit for their contributions to science. I provide no explanation for these characteristics of scientists. The human species seems innately inquisitive. The process by which young people are introduced into science at least does not totally destroy this native curiosity. In some cases, it even encourages it. According to Harré (1979), the desire for recognition from one's peers is equally strong in human beings. Even if a budding young scientist enters science not caring about something as paltry as individual credit, he or she will find it very difficult not to get caught up in the general enthusiasm. I have also not presented any justification for our belief in an external world which we can come to know or for the existence of any regularities in nature. However, given curiosity, a desire for credit, and the possibility of checking, the structure that I claim characterizes science can explain quite a bit about the way in which scientists behave.

Many commentators find one or more features of science as it has existed for the past couple hundred years to be less than palatable. They find scientists too polemical, aggressive, arrogant, and elitist. Scientists are too anxious to publish so that they can scoop their competitors. They seem more intent on enhancing their own reputations than in helping humanity. According to these commentators, scientists should turn their attention from the problems that they find most interesting to those that are currently most relevant to our survival. The mechanism which I propose explains why scientists do not behave in the way that these critics think that they should behave. From an operational point of view, behavioral psychologists are right on at least one count: organisms tend to do what they are rewarded for doing, pious hortatory harangues notwithstanding. If scientists are rewarded for making new discoveries, formulating more powerful theories, designing novel experiments, etc., then they are likely to do just that. Perhaps scientists could be raised so that they were not so strongly motivated by curiosity and the desire for individual credit, but I am not sure that the results would be worth the effort. In fact, such efforts, if successful, might bring science to a halt. 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.

The mechanism that I have sketched in this paper may not seem like much of a mechanism, hardly up to explaining the marvelous progress made during the past few centuries by successive generations of scientists. But when one thinks of it, natural selection is not much of a mechanism either, and yet it has produced all the fantastic adaptations that organisms, both extinct and extant, exhibit. The mechanism that I sketch is also not very efficient. If nearly all the progress in science turns on the work of a very small percentage of scientists working at any one time, then science could be made much less expensive and no less efficient by weeding out those scientists who are not being very effective. However, if biological evolution has any lessons to teach us, it is that selection processes cannot be made too efficient without neutralizing their effects. The sort of interindividual and interdemic polemics that have characterized science from the beginning are not a very efficient way of reconciling differences among competing scientists and groups of scientists, but they are extremely effective.

NOTES

* I wish to thank both Michael Ruse and Ronald Giere for suggesting improvements in an early draft of this paper. This paper is an abstract of a very long book. The number of people who helped me in developing the ideas set out in this book is extremely large, so large that I decided to defer expressing my gratitude to them until its appearance.

[1] In addition to omitting data from this paper, I have kept references to a minimum. Both of these omissions are serious, given the mechanism for conceptual change that I set out in this paper.

[2] From past experience, I have learned that the preceding distinctions are difficult to keep straight, in part because ordinary English is not constructed to make them. I am arguing that the species category might well be a natural kind with a central role in biological evolution, that *Homo sapiens* might well be an instance of such a natural kind, but that the human species itself is not a natural kind. Parallel distinctions with respect to a physical element such as gold are as follows: physical element is a natural kind. One instance of a physical element is gold. It too is a natural kind. However, the pope's ring is only an instance of gold.

[3] "Vehicle" is used in two distinct senses in the literature on selection processes. Campbell (1979) uses it as I do to refer to the physical entities that are functioning as replicators. Dawkins (1976) uses it to refer to interactors. Williams (1985) prefers not to term organisms "vehicles" because it plays down the active role of organisms in biological evolution. Needless to say, I prefer Williams' usage.