

DUBLIN CITY  
UNIVERSITY

Ollscoil Chathair Bhaile Átha Cliath

Essays on Darwinism  
1: Ontological Foundations

Barry McMullin

March 1992

(Minor corrections: 14 April 1992)

**School of  
Electronic Engineering**

TECHNICAL REPORT: bmcm9201

©1992 Barry McMullin,  
School Of Electronic Engineering,  
Dublin City University,  
Dublin 9,  
IRELAND

**Telephone:** +353-1-704 5432  
**Fax:** +353-1-704 5508  
**E-mail:** McMullinB@DCU.IE

## Abstract

The problem tackled in this essay is to formulate “Darwinian” theory in a way which is divorced from any specifically biological interpretation. For my purposes, the crucial constraint on this abstract formulation of Darwinian theory is that it must qualify as a realisation of *unjustified variation and selective retention* in the sense of D.T. Campbell (1960a; 1960b; 1974a; 1974b)—for it is (arguably) only thus that Darwinian evolution can lead to the spontaneous growth of “knowledge”. I hope that such an abstract formulation might help to clarify the application of the theory even in its native biological setting; but, more importantly, it seems to me that this exercise is a necessary pre-requisite for the valid application of Darwinian principles in any domain *other* than conventional, terrestrial, biology—such as the emerging field of study commonly referred to as *Artificial Life*. In attempting to solve the problem of formulating a satisfactory, abstract, version of Darwinism, I introduce a novel ontological category: the Similarity-lineage, or *S-lineage*. With the aid of this hybrid category, which combines some aspects of both a class and an individual, I argue that Darwinian evolution relies on a process of S-lineage selection—which is to say that S-lineages are, in a certain precise sense, the elusive Darwinian *units of selection*.

# Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
<b>2</b>	<b>Metaphysics</b>	<b>2</b>
2.1	Individuals . . . . .	2
2.2	Classes . . . . .	2
<b>3</b>	<b>Actors</b>	<b>3</b>
<b>4</b>	<b>Lineages</b>	<b>4</b>
<b>5</b>	<b>On Similarity</b>	<b>5</b>
<b>6</b>	<b>Selection</b>	<b>6</b>
6.1	Determinism (sort of) . . . . .	6
6.2	Malthus . . . . .	7
6.3	Competition . . . . .	8
6.4	Consistency . . . . .	8
6.5	S-creation . . . . .	9
6.6	Natural or Artificial? . . . . .	12
6.7	In Summary . . . . .	12
<b>7</b>	<b>Darwinian Evolution</b>	<b>13</b>
<b>8</b>	<b>Climb Every Mountain</b>	<b>14</b>
<b>9</b>	<b>On Adaptation</b>	<b>15</b>
<b>10</b>	<b>M.B. Williams: Darwinian Axiomatics</b>	<b>17</b>
<b>11</b>	<b>The Unit of Selection</b>	<b>19</b>
<b>12</b>	<b>Conclusion</b>	<b>21</b>
	<b>References</b>	<b>22</b>

# 1 Introduction

With the work of Einstein, physicists were faced with a dilemma: If they insisted on retaining Euclidian geometry, they would have to be content with extremely complicated and variable laws; if they wanted laws applicable anywhere in the universe, regardless of velocity, they would have to abandon Euclidean geometry. They opted for the second alternative. Evolutionary biologists are currently confronted by a similar dilemma: If they insist on formulating evolutionary theory in terms of commonsense entities, the resulting laws are likely to remain extremely variable and complicated; if they want simple laws, equally applicable to all entities of a particular sort, they must abandon their traditional ontology.

Hull (1980)

The general notion of “Darwinian evolution” may be used very loosely, with a wide variety of interpretations. Considerable confusion can result. Thus I believe that it may be useful to present a general, and abstract, discussion of the “concept” of Darwinism, and that is the primary purpose of this essay.

This should not be interpreted as an attempt to probe the “true essence” of Darwinism but rather as a stipulation of the kind of thing to which I have decided to attach the label “Darwinism”. As it turns out, my usage is somewhat distinct from, or at least more precise than, the more generally accepted usages, though naturally it overlaps considerably with them. I should emphasize that this increased precision is not pursued for its own sake, but because I have found it necessary for the treatment of problems of interest. Nonetheless, some readers may wish to challenge my choice of “Darwinism” to denote this thing; in which case, I should refer them to Karl Popper:

If challenged by the question of whether a word one uses really means this or perhaps that, then one should say: ‘I don’t know, and I am not interested in meanings; and if you wish, I will gladly accept *your* terminology.’

Popper (1979, Essay 8, p. 310)

We might expect that Darwinism could most easily be characterised simply by reviewing it in its original, biological, context. However, I am in agreement with the general position outlined in the open-

ing quotation above from Hull. The orthodox categories used in biology predate Darwinism; thus it would be hardly surprising if they proved not to be the most suitable vehicles for the expression of Darwinian theory, even within biology itself. The result is that, faced with the complexity of modern evolutionary biology, it has become very difficult to distinguish issues which are intrinsic to the notion of Darwinism from details or artefacts arising from the particular realisation(s) of Darwinian principles which (happen to) exist in the biological world as we know it.<sup>1</sup> I would even go so far as to say that the vastly more comprehensive biological knowledge we now have, relative to that available to Darwin, may actually serve to obscure rather than illuminate his original insight.

The presentation I give here has been substantially influenced by the ideas of Hull (1980; 1981), and Dawkins (1976; 1982a). Nonetheless, I present an essentially *original* proposal (complete with its own jargon). I take this step with some reluctance, but no apology. I shall argue elsewhere (McMullin Forthcoming) that the analyses of Hull and Dawkins harbour some subtle and intricate flaws; but it seems to me that *that* argument cannot even be coherently stated without first establishing an independent ontological foundation to support it.

Quite separately, there is also a significant overlap between my formulation and the *Axiomatic Darwinism* introduced by M.B. Williams (1970) (although I became aware of this connection only after the essay was largely complete). I review this in more detail in section 10 below.

The important issue in the framework I present is, of course, its substance—the ontological categories I introduce, and the relations between them; the specific terminology is secondary. However, I should warn in advance that at least some of my terminology may seem somewhat arcane. My only excuse is that it was necessary to carefully discriminate all the required concepts while still avoiding any existing terms which might have had confusing or prejudicial connotations. This is not easy. In any case, I can do no less than accept the same assessment criteria as Hull suggested for his own related work:

This reconceptualization of the evolutionary process is certainly counter-intuitive; its only justification is the increased scope, consistency, and power of the theory that results. If the terminology suggested in this paper cannot characterize the evolutionary process more accurately and suc-

---

<sup>1</sup>Implicitly, of course, I conjecture that some such distinction is indeed possible.

cinctly than the traditional terminology, it should not (and will not) be adopted.

Hull (1980)

## 2 Metaphysics

I shall use the following general metaphysical categories. These are based on categories proposed by Hull (1981), but with some qualifications.

### 2.1 Individuals

In Hull's words, an individual is a "spatiotemporally localised entity which develops continuously through time". That is to say, it is an entity which occupies a connected region of spacetime—it has a continuous (but generally finite) extension in space and time. Implicitly, I am accepting "space" and "time" as ontological primitives requiring no further analysis.

Individuals must be reasonably discrete in both time and space. In particular, they can (if necessary) be distinguished from each other *solely* by reference to their extension in time and/or space—no two individuals may have precisely the same spatiotemporal extension. That is, individuals are unique historical entities. I shall refer to the extension of an individual in time as its *lifetime*, and say that this lifetime is delimited by its *birth* and its *death*.

Hull adds that individuals should also exhibit "internal cohesiveness at any one time". I presume he introduces this condition in order that completely arbitrary regions of spacetime need not be recognised as individuals. However, I see no benefit in making this *critical* for individuality—in practice we are not so capricious in our designation of individuals. If an entity's *only* characteristic is its extension in spacetime, we shall not be interested in it anyway; otherwise, we may well be interested in it, and argument about whether its characteristics include "cohesiveness" (whatever that may be) are hardly productive. Indeed, Hull himself makes a similar point when he says that "the only individuals which are scientifically important are those which are acted upon by natural processes". No more stringent criterion is needed.<sup>2</sup>

Distinct individuals, although unique, may be said to be more or less *similar* according to various

---

<sup>2</sup>If, notwithstanding my arguments here, one insists on including "cohesiveness" in the definition of "individual" I suggest it can be best done by interpreting it in the relatively technical sense of *autopoiesis* (Varela 1979). I am, in fact, quite sympathetic to that view; I just feel that it would be an unnecessary complication for my immediate purposes here.

criteria. In particular, they will be said to be "identical" if they differ *only* in spatiotemporal location. This is, of course, a purely metaphysical notion of identity—I do not imagine that it could be operationally realised—and for that reason I shall retain the scare quotes around it. I take Hull to have the converse of this concept in mind when he refers to the idea of individuals being "numerically" distinct.

Note that we have here a specialised, technical, definition of "individual". In particular, it must not be confused with the common biological convention which uses "individual" almost synonymously with "organism".

### 2.2 Classes

A class is a *spatiotemporally unrestricted* entity, which can have individuals as members. Thus, although a class consists of, or is composed of, entities which are spatiotemporally localised, the class itself is not—it exists outside the spacetime framework. Indeed, a class may be perfectly well-defined even if it has no members.

For many purposes, classes can be thought of as mathematical *sets* (with the stipulation that class membership must be spatiotemporally unrestricted). In what follows I shall freely, if loosely, use the concepts and terminology of set theory in discussing classes.

Class membership cannot involve spatiotemporal relations between individuals, as this would imply some kind of spatiotemporal restriction (relative or absolute) on the class.

The definition of a particular class *introduces* a relation between its members, of course; but this relation will not be recognisable from, or dependent upon, the spatiotemporal configuration of the individuals. To put it another way, the class membership of an individual must not be sensitive to its spatiotemporal location. In particular, given that one individual is a member of a class, it follows that all "identical" individuals (if any), regardless of spatiotemporal location, must be (distinct) members of the class. These need not, of course, be the *only* members of the class—i.e. classes are not restricted to consisting of "identical" individuals.

Apart from the requirement that it not involve any spatiotemporal restriction, no other constraint is placed on mechanisms of class definition—although, obviously, for any practical application of the ideas discussed here, we would require that class membership be defined *operationally*.

Classes, being spatiotemporally unrestricted, are evidently not, themselves, individuals. So classes

and individuals are mutually exclusive kinds of entity. In what follows, an entity may be described as a class or an individual, but never both.

### 3 Actors

Individuals may be related in various ways. Of particular interest for my immediate purposes are relations of *descent*. Hull relies on the common sense idea of descent: that individuals are related by descent when one entity “causally produces” another. Without disagreeing with Hull on this, I should nonetheless like to define the notion a little more carefully, as I shall ultimately want to consider certain distinct kinds of relation as being examples of (distinct kinds of) descent. It is important at this point that the reader try to put aside (for the time being) any preconceptions she may have as to what exactly constitutes a relation of “descent”.

I shall refer to an individual for which a relation of descent is defined as an *actor*. It is important to emphasize that individuals are not intrinsically or self-evidently actors—they are actors only relative to some definition or interpretation of what we mean by descent. Indeed, the same individual may be an actor relative to more than one definition of descent. So whenever an individual is referred to as an actor below, some particular definition of descent, applicable to that individual, is always being assumed, even if this is not expressly stated.

While I have said that descent is a matter of definition, I do not mean that any arbitrary relation defined on some class will serve as a relation of descent. There are certain restrictions which it must conform to.

Descent in general is defined in terms of a more primitive relation, which I shall call *immediate descent*. This relation has the following characteristics:

- It is *binary*, i.e. it is a relation which exists (or not) between *two* individuals.
- It is asymmetric—it has a “direction”. I shall henceforth speak of one individual being (immediately) descended *from* another; the former will be termed the *offspring* and the latter the *parent*.
- In general, any particular definition of immediate descent is applicable only to some particular set of individuals. I shall assume that such a set will be spatiotemporally unrestricted—i.e. it will constitute a class. Thus, a definition of immediate descent will incorporate, implicitly or

explicitly, the definition of a class of individuals for whom this relation can be evaluated. This class is then, by definition, the class of *actors* relative to the given definition of (immediate) descent.

- I shall stipulate that an actor can be the parent of another only if the birth of the second occurs *after* the birth of the first. Informally of course, we mean that, in some sense, the parent has “causally produced” the offspring, or at least has been “causally implicated” in its production. It follows that no actor may be a parent of itself. This condition is somewhat similar to M.B. Williams’ (1970) *Axiom B1*; although she avoids any explicit spatiotemporal restriction, the effect is essentially equivalent. My formulation is, in one sense, slightly stronger in that we can already also deduce that no actor can be a parent of (any of) its own parent(s), whereas Williams provides a separate axiom (*B2*) to capture this. On the other hand, Williams’ formulation has the advantage that it does not require an axiomatic conception of absolute time.

Now a definition of immediate descent can, and generally does, depend on a variety of factors relating to the particular actors; but regardless of what else it depends on, it is clear that if one actor is immediately descended from another this fact necessarily implies something about their relative spatiotemporal locations. Thus, no set of individuals which is defined wholly or partly in terms of immediate descent can function as the definition of a class.<sup>3</sup>

Immediate descent requires “similarity” between parent and offspring only in the following weak sense: they must both (by definition) be members of the class on which this particular form of immediate descent has been defined. Thus immediate descent must, in general, imply *some* kind of similarity between parent and offspring (namely, enough to establish that they *are* both actors relative to the given definition). To put it another way, immediate descent may be compatible with more or less arbitrary *differences* between parent and offspring. It is for this reason—the fact that immediate descent, as

<sup>3</sup>At the risk of compounding the confusion already engendered: a *definition* of immediate descent may be spatiotemporally unrestricted, and may serve to define a class (the class of individuals on which this definition of descent may be successfully evaluated—its actors); however, any particular *example* of immediate descent, in accordance with some such definition, necessarily contains spatiotemporal restrictions and cannot function wholly or partly as the definition of a class.

I use it, does not *necessarily* require a high “degree” of “similarity” (however that might be defined) between parent and offspring—that I do not use terms like “replication” or “reproduction” here.

I have expressed (or at least, constrained) immediate descent as an asymmetric, binary, relation between individuals (actors relative to this form of immediate descent). That is, any two actors definitely either are or are not related by immediate descent; and if they are so related, one is definitely identifiable as the parent and the other as the offspring. This should not be read as restricting the number of such relations any single actor may participate in: in general, an actor may have zero, one, or many parents, and may have zero, one, or many offspring. Particular definitions of immediate descent may be more restrictive of course.

If an actor has no parents, I shall say that it has been *created*; otherwise I shall say it has been *pro-created*. A created actor will also be called a primary actor, or simply a *primary*.<sup>4</sup> Please note carefully that I use the term *creation* to describe the birth of an actor, *other* than by (some particular definition of) immediate descent. *This is without prejudice to the exact process actually involved*. For example, creation relative to one definition of descent may actually correspond to *pro-creation* by some other definition. Thus when I refer to creation I do not imply any inherently supernatural or mystical process. I am aware of the dangers of my usage, which already existed, perhaps even more acutely, in Darwin’s day. As Hull puts it:

...even though ‘creation’ might well have been a code word for unknown natural processes, the use of this word, especially when it was coupled with all sorts of additional theistic references, was guaranteed to give just the opposite impression.

Hull (1983)

However, it seems to me that “creation” is still, in its literal sense, the best word for the job; so, I will continue to use it, while hoping that this advance warning will mitigate any confusion.

If an actor has exactly one parent, I shall say it is a product of *unimodal* procreation; if it has exactly two parents, I shall say *bimodal* procreation etc.; if the number of parents is unspecified, but definitely more than one, I shall refer to *multimodal* procreation. This is obviously similar to the biological idea of asexual versus sexual reproduction, but is more

---

<sup>4</sup>Primaries thus constitute a sub-class of the class of actors, relative to some particular definition of (immediate) descent.

general in that it permits more than two parents, and is neutral on the question of whether there are any constraints on which actors may “mate” with each other—in particular, whether there is any kind of differentiation of actors into distinct and complementary “sexes”.

A set of actors which are all related by immediate descent to a certain parent are said to be *siblings*.

Given a relation of *immediate* descent, the generalised relation of descent is simply the transitive form of this. That is, if one actor is a parent of a second, and it, in turn, is a parent of a third, then the first is said to be an *ancestor* of the third, and the third is said to be a *descendant* of the first (and so on).

## 4 Lineages

Given a relation of descent, I shall define a *lineage* as any set consisting of some identified *primary* actor (called the *founder* of the lineage), together with all of its descendants.<sup>5</sup> Thus, each actor in a lineage is connected to the founder by at least one continuous chain of ancestors each of which is also a member of the lineage.

Every lineage has exactly one founder, and every primary actor founds exactly one lineage (possibly consisting only of itself). Any given actor may be a member of an arbitrary number of distinct lineages—specifically it will be a member of the lineages founded by each of its primary ancestors. However, by the same token, a primary actor is a member of just one lineage, being the one it founds.

Under any form of unimodal (“asexual”) procreation every actor is either itself a primary or is descended from one, unique, primary ancestor. Thus, in this case, lineages will all be disjoint (every actor will be a member of exactly one lineage). However, where multimodal (“sexual”) procreation is possible, lineages may, over time, converge to an arbitrary extent. Note that if it arises at any time that the membership of two lineages has become equal, then the membership of the two lineages will remain identical thereafter—effectively they have become one.

Lineages are, in the terms discussed previously, well-defined *individuals*; that is, they are distinct

---

<sup>5</sup>I diverge somewhat from Hull’s (1981) definition of “lineage” here: Hull does *not* require the founder of a lineage to be a primary (indeed, in Hull’s terms, *every* actor may be regarded as a founder). However, I will effectively recapture the required generality of Hull’s version when I introduce the notion of an *S-lineage* in section 5 below. I may note in passing that the Hullean lineage is, in turn, largely equivalent to M.B. Williams’ terms *clan* and *subclan* (Williams 1970).

entities which are reasonably well localised in space and time (by virtue of consisting of a definite set of actors, each of which is separately well localised in space and time). Although it is an individual, a lineage is, of course, a distinct individual from any of its members; indeed a lineage is, in general, a different *kind* of entity from its members. In particular, a lineage will not be an actor—not, at least, with respect to the same definition of descent as relates its own members. A lineage may, of course, function as an actor relative to some (other) definition of descent—in which case, there will exist “higher level” lineages etc.

Subsets, intersections and unions of lineages are not generally lineages in their own right (though they are always well-defined individuals).

Although a lineage has been defined as a kind of set, the “size” of a lineage will not refer to its total number of members (or its cardinality in conventional set theoretic terms), but rather to the number of members *at any given time*. In this sense, lineages may grow or shrink over time and this will be reflected in their sizes, which will be, in general, well-defined integer functions of time.

A lineage continues in existence for as long as some descendants of the founder have further offspring, so that a lineage is *potentially* immortal (even if the actors themselves necessarily each have a limited lifetime). However, a lineage can and does die if it happens that, at any time, the last remaining member dies—which is to say that the size of the lineage becomes zero.

## 5 On Similarity

I have been at some pains to emphasize that a relation of descent, in itself, need say very little (if anything) about the “similarity” between parents and offspring. We shall see that this freedom is essential: Darwinism demands that *both* similarity *and* variation be associated with descent in certain circumstances.

Let there be a class of individuals which are actors relative to some relation of descent. Let an arbitrary, but well-defined, sub-class of this class of actors be termed a Similarity-class or *S-class*. That is, S-classification is a formalisation of the notion of “similarity” *within* some class of actors. Note that any given actor can, in general, be a member of many different S-classes.

Consider some actor which is a member of a particular S-class. We shall say that if an offspring of this actor is *also* a member of this S-class then this is an example of *S-descent* (relative to the specified S-

class). S-descent is a special case of ordinary descent in which the offspring is “similar” to the parent in whatever characteristic(s) determine membership of the particular S-class.

Depending on our choice of S-class, S-descent may range from being non-existent to universal. Roughly speaking, the former would correspond to an S-classification which is not heritable at all (is independent of the given relation of descent), and the latter would correspond to an S-classification which is completely heritable (in the limit, the S-class is set equal to the class of actors for the given relation of descent, which classification is completely “heritable” by definition).

A Similarity-lineage or *S-lineage* (relative to some S-class) is then defined in much the same way as an ordinary lineage, except that the members must be related by *S*-descent. More strictly, we define an S-primary as a member of the specified S-class, none of whose parents (if any) were in the S-class; we define an S-descendant of an S-primary as any actor connected to the S-primary by a continuous chain of S-descent; and an S-lineage is then defined as a set consisting of an S-primary together with all its S-descendants. Thus, every S-primary is an S-founder of an S-lineage. The S-size of an S-lineage will refer to the number of members it has at any particular time. The notion of an S-lineage is largely equivalent to Williams’ (1970) *Darwinian subcland* (*sic*).

S-lineages will evidently be subsets of lineages, but will *not* be lineages in their own right (an S-primary need not be a primary in the absolute sense; for that matter, an S-lineage will, in general, exclude some descendants of the S-founder—namely all those which lack a continuous chain of *S*-descent from the S-founder).

By the definition of S-descent, we can see that all actors making up a single S-lineage must be members of the same S-class. This is the peculiar manner in which the S-lineage idea is a sort of hybrid of an individual and a class: it is (among other things) a spatiotemporally delimited subset of a class. In what follows I shall sometimes loosely speak of an S-lineage being “of” an S-class as a shorthand for saying that all its *actors* are of that S-class.

Note, however, that the converse to this idea is not true: actors which are of the same S-class need not be members of the same S-lineage. Specifically, if there exists more than one S-primary belonging to a given S-class, these will found multiple S-lineages containing actors of the same S-class.

In the case of unimodal (“asexual”) procreation, all S-lineages relative to a given S-class, being subsets of distinct lineages proper, are necessarily dis-



joint. In the case of multimodal (“sexual”) procreation, on the other hand, S-lineages relative to the same S-class may intersect; indeed, at some point in time they may become completely converged, and must then remain completely converged at all later times.

S-lineages relative to *different* S-classes may intersect under both unimodal and multimodal procreation, provided only that the S-classes themselves intersect—which is to say that a single actor may be a member of many different S-lineages of different S-classes; or conversely, if two S-classes are necessarily mutually exclusive, then their respective S-lineages (if any) cannot intersect. Even if two S-lineages relative to different S-classes become completely converged at some time, they may subsequently diverge again.

The actors present in a system at any given time may be viewed as distributed over one or more S-lineages relative to various S-classes. These S-lineages may grow and shrink over time, and, on occasion, some may “die” (become extinct). To the extent that there is an on-going process of *S-creation* (the birth of new S-primaries) new S-lineages may also be initiated or “born” over time.

## 6 Selection

I shall now consider the general notion of *selection*, as it arises in Darwinian theory. For reasons which will become apparent, I shall refer to this as *S-lineage selection*. Loosely speaking, S-lineage selection is a process in which one S-lineage consistently grows, at the expense of another, ultimately displacing it.<sup>6</sup>

The purpose of this section is to try to characterise as precisely as possible the circumstances under which such a process will actually occur. I shall concentrate on the relatively simple case of selection between just two S-lineages: the generalisation to more than two should be reasonably clear, and will not be discussed explicitly.

### 6.1 Determinism (sort of)

Our first requirement may be informally stated as being that the changes in the S-sizes of each of the S-lineage’s in question should be “determined” or

---

<sup>6</sup>In this respect, my usage of “selection” is more restrictive than is usual in population genetics—where the word seems to be commonly used to mean *any* “approximately deterministic” change in S-size. The latter condition is detailed in section 6.1, but as only one among a number of conditions which are required for “selection” in my sense.

“predicted” by their S-classes. This is a difficult idea to pin down precisely, but I shall make the attempt nonetheless.

Clearly, the changes in S-size may generally depend on a variety of factors “outside” of the S-lineage of interest. These include the effects of actors not in the S-lineage, and any salient aspects of the common system in which the actors are all embedded. These things will be collectively referred to as the *S-environment* (each S-lineage clearly defines its own unique S-environment). We can thus restate our requirement, somewhat more precisely, as being that S-class should be predictive of the rate of change of S-size (per unit time) for any *given* S-environment.

It is important to see that this requirement is non-trivial; it certainly need not hold for arbitrary S-classes. In principle, even if all “relevant” aspects of the S-environment have been stipulated, the S-class of the actors making up the S-lineage may still not allow any prediction of the rate of change in S-size: the characteristics upon which the S-classification are based may bear no relationship whatsoever to rate of change in S-size. More generally, these characteristics may have *some* effect on the rate of change in S-size, but this effect may be swamped by other characteristics which are not relevant to the S-classification. That is, two distinct S-lineages (relative to the same S-class), subjected to the same S-environment, may still exhibit radically different rates of change in S-size, simply because the particular S-class does not capture other *relevant* characteristics of the actors.

Having recognised that *arbitrary* S-class need not be predictive of the rate of change in S-size, we can now ask what it might mean to say that a particular S-class *is* predictive in this way. I shall try to answer by first giving a somewhat contrived and restrictive example, and then indicating, in outline at least, how it might be generalised.

Suppose then that the S-class of an actor (together with the S-environment) uniquely determines two probability functions: one specifies the probability that the actor will die in a unit time interval, and the other specified the probabilities that the actor will have each possible number of S-offspring in a unit time interval.<sup>7</sup> We stipulate that these probability functions are stationary—i.e. do not vary with time. Now, in this case, S-class is clearly implying

---

<sup>7</sup>In the case of multimodal (“sexual”) procreation, we would require some convention such that a given S-offspring is counted only relative to one S-parent, so that we can subsequently infer information about the total change in S-size. But this is a minor detail not affecting the general principles here.

very strong and particular similarities between actors of that class, which was not the case for arbitrary S-class. It also implies that actors of this S-class do not interact with each other—or, at least, do not interact in any way which affects these probability functions (the functions are independent for every actor).

It then follows that there will exist two derived probability functions, which will be associated with any S-lineage relative to this S-class. The first is the probability function for the number of actors in the S-lineage which will die per unit time, and the second is the probability function for the number of S-offspring which will be born per unit time. Both of these will depend on the current S-size. I shall term the death *rate*—which is to say the number of deaths per unit time, divided by S-size—the *S-mortality*. Similarly, the birth rate will be termed *S-fecundity*. S-mortality and S-fecundity will be random variables. Under the conditions described, their expected values will be independent of S-size (they will, in fact, be equal to the expected values of the corresponding probability functions for the individual actors). Furthermore, their variances will be inversely related to S-size.

That is, S-class (together with S-environment) will predict the probability functions for S-mortality and S-fecundity; and, further, if S-size is “sufficiently” large, S-mortality and S-fecundity will, in effect, be uniquely determined by S-class. The rate of change of S-size is simply the difference between S-fecundity and S-mortality (multiplied by current S-size).

This example shows the strongest case of what might be meant by S-class being predictive of the rate of change in S-size. I now want to suggest, even if only in outline, how this case might be generalised.

The essential extra idea is to allow for diversity or structure *within* an S-lineage. That is, we don’t require that every actor of the given S-class should have exactly the same probability functions for death and numbers of S-offspring. Instead we allow that different actors may have different probability functions; but we insist that the distribution of these distinct probability functions *should* be determined by S-class. Crudely speaking, we might envisage that the actors are organised into “social” units, such that these social units would all exhibit the same gross probability functions for births and deaths, even though there is heterogeneity *within* each social unit. Provided the S-size of the S-lineage is large enough to contain many such units we will still have the situation that S-class is more or less deterministically predictive of the rate of change of

S-size.

Note the implication here that, *in general*, it is quite misleading to think of S-fecundity and S-mortality as directly reflecting “corresponding” characteristics of the constituent actors. For example, while any given actor can contribute to S-fecundity by the obvious mechanism of bearing S-offspring itself, it can also significantly influence S-fecundity by helping (or hindering) other members of the S-lineage in bearing offspring. That is:

*S-fecundity and S-mortality are characteristics of S-lineages, not actors.*

In summary then, our first requirement for S-lineage selection is that S-fecundity and S-mortality of an S-lineage should be uniquely determined (at least “probabilistically”—in the sense of definite, stationary, probability distributions) by its S-class and some specified aspects of its S-environment. This allows S-fecundity and S-mortality to have more or less arbitrary dependencies on the S-environment; what it rules out (indeed, almost the only thing it rules out) is the possibility that S-fecundity and S-mortality depend on any aspects of the S-lineage itself *other* than those implicit in the S-class of its actors.

The only remaining variability in S-fecundity and S-mortality (if any) is limited to probabilistic variation according to a stationary probability function. Which is to say that, provided the size of the S-lineage is large enough, then (by appeal to the so-called “law of large numbers”) this remaining variability can be made arbitrarily small. I do not, of course, attempt to give any particular figure for how large would be large “enough”—that will depend on the particular case; it is sufficient if, in principle, some such figure can always be specified.

Technically, the requirement which has been stated is that, provided the size of the S-lineage (denote it  $s$ ) exceeds some minimum (i.e. is large “enough”) then it will satisfy, to an arbitrarily close approximation, a deterministic differential equation of the form:

$$\frac{ds}{dt} = (F - M)s$$

where  $F$  and  $M$  denote S-fecundity and S-mortality respectively, and are uniquely determined by S-class, in the given S-environment.

## 6.2 Malthus

My next requirement for S-lineage selection is that the S-lineages in question be “capable” of increasing more or less exponentially in size provided sufficient raw materials or resources are available. This

is the so-called “Malthusian” growth that originally inspired both Darwin and Wallace.

I have already stipulated that S-class, in combination with S-environment, should determine S-fecundity and S-mortality (at least if S-size is large “enough”). I now require that this dependence on S-environment be expressed in terms of “resources”; and that provided these resources are sufficiently “large”, it must be possible for S-fecundity to exceed S-mortality ( $F > M$  in the previous equation). S-size will then grow (exponentially in the first instance). As S-size increases however, and resources<sup>8</sup> therefore become more scarce, S-fecundity will fall and/or S-mortality will increase. Eventually this will result in a balance between the resource availability and S-size, such that S-fecundity and S-mortality are just equal, and S-size remains more or less constant.

### 6.3 Competition

Our third condition for S-lineage selection is that the two S-lineages must *compete* for certain limited resources—this is what Darwin (1859, Chapter III) termed the “struggle for existence”. This is actually quite subtle and the requirements must be spelled out carefully.<sup>9</sup>

We require two things. Firstly, the two S-lineages must both rely on the same resource(s). Thus, an increase in the S-size of either will have an adverse effect on the resource availability of both. This is a relatively straightforward requirement.

Secondly, and less obviously, we require that the two lineages are, to at least some extent, *exclusive* of each other. More technically, if  $s_1$  and  $s_2$  are the respective S-sizes, we require that:

$$\frac{d s_1}{d s_2} < 0$$

This requirement is trivially satisfied if the S-lineages in question are disjoint: in that case an increase in one can, in the face of fixed resource availability, occur only at the expense of a reduction in the other. In the case where the S-lineages are not disjoint (i.e. under multimodal procreation,

---

<sup>8</sup>Strictly, this should be interpreted in terms of fixed resource “flows” (per unit time).

<sup>9</sup>Indeed, at first sight it seems that Lewontin (1970, p. 1), in particular, holds that resource limitation is *not* a necessary condition for selection; but I think this is because he uses “selection” in that generalised sense, noted earlier, of *any* (approximately) deterministic change in S-size, regardless of whether this involves the eventual elimination of any S-lineage. By contrast, I reserve “selection” to denote cases where some S-lineage(s) really are selected in the sense of displacing or eliminating some alternative(s).

and assuming that the S-class definitions are not “inherently” disjoint), this condition evidently may or may not be satisfied.

### 6.4 Consistency

Our third requirement for S-lineage selection is that there should be a “consistent” bias in the changes in S-size, in favour of one or the other.

That is, our previous requirements have established that the changes in the S-sizes of the two S-lineages will be more or less deterministic and must, instantaneously, be in opposite directions. However, that would still allow a situation in which the S-sizes could stabilize (both rates of change become zero) or could oscillate (for a while, one increases and the other reduces, and then vice versa etc.).<sup>10</sup>

This implies two distinct things in turn. Firstly, we must require that the S-environments of the two S-lineages do not change “too much”. That is, we can permit some variation in the S-environments provided that such changes affect both S-lineages “equally” (or, at least, in the same direction). We have already explicitly introduced a change in S-environments of this kind in the form of the dependence on shared “resources”. The point now is to restrict changes in the S-environments to be *only* of this general kind (at least for the duration of the selection “episode”).

Secondly, we must assume that whatever “bias” exists must be consistent for arbitrary S-sizes of the two S-lineages. By “bias” here I mean, roughly, the difference in the two rates of change (per actor in the S-lineage). If the first S-lineage increases when its size is small compared to the second, it should still increase even if (when) it becomes large compared to the second. I do not require that the bias should remain strictly constant (be independent of the S-sizes), but just that it should stay in the same *direction*.

It now “almost” follows that one S-lineage will, over time, necessarily completely displace or eliminate the other; which is to say that one of the competing S-lineages will be *selected*. I say “almost” because there is still one further condition to be satisfied—but I shall return to that in the next section.

Selection bias is, in some sense, a measure of the rate at which selection is instantaneously taking place. I have allowed that it may vary; but the

---

<sup>10</sup>Examples of this kind of phenomenon can arise through so-called *frequency dependent selection*—which is again something of a misnomer in my terms, as it implies, *inter alia*, that no single S-lineage need be actually *selected*.

“important” aspects of it (its direction, and approximate magnitude) may be indicated by giving its value under some kind of “standardised” conditions. Quite what standardised conditions are most convenient depends on the detailed dynamics of the particular system, and I shall not go into that. I merely wish to note that it may be possible to attribute standardised numbers to particular S-lineages (in standardised S-environments) in such a way that, for any pair of competing S-lineages, the difference in these numbers will serve to predict the outcome of selection between them (which will displace the other, and how quickly). I shall refer to such standardised numbers as *S-values*.

The notion of S-value is closely related to the concept of *fitness* as it is commonly used in population genetics. However I deliberately avoid the use of this latter word because it has been used in a variety of distinct, sometimes mutually contradictory, ways—even within Darwinian theory. This problem has been documented in detail by Dawkins (1982a, Chapter 10), where he, too, concludes that the word has become so overloaded that it is now best avoided altogether, if at all possible.<sup>11</sup>

## 6.5 S-creation

The final condition for S-lineage selection is that it should not be “compromised” by *S-creation* events.

S-creation means the birth of an actor of a given S-class, none of whose parent(s) were of that S-class. S-creation results in the foundation of a new S-lineage (it would typically correspond to the occurrence of a “mutation” in conventional biological terms).

I consider three cases here.

Firstly, S-creation could result in the foundation of an S-lineage of some S-class which is different from the S-classes of either of the pre-existing S-lineages, *and which does not compete with them*. If effect, this corresponds to the foundation of an S-lineage which exploits a distinct environmental “niche” (which may, or may not, mean that it must compete with some S-lineages *other* than those originally under consideration). This may, or may not, have some effect on the selection dynamics between the original two S-lineages; but if it does have such an effect, this falls under the heading of a significant “environmental perturbation”, and we would no longer be dealing with S-lineage selection *per se*.

<sup>11</sup>My “S-value” is also essentially the same as the “selective value” described by Provine (1986, Chapter 9)—which is something of a happy coincidence: the “S” in “S-value” strictly stands for “similarity” (being derived from “S(imilarity)-lineage”), rather than for “selection”.

This general case is very complex, and I shall not pursue it further here; I shall return to it, briefly, in section 8 below.

Secondly, S-creation could result in the foundation of an S-lineage of some S-class different from the S-classes of either of the pre-existing S-lineages, but which *does* also compete with them (i.e. meets all the conditions already identified for this). Clearly, this new S-lineage will then potentially interfere with the completion of the selection episode between the original two S-lineages. On the other hand, it simply means that we must generalise our concept of selection to allow selection from more than two candidates—it does not fundamentally affect our arguments so far, and I shall not consider it further.

The final case is where S-creation results in the foundation of an S-lineage of the *same* S-class as one or the other of the pre-existing S-lineages. This possibility introduces some subtle, and fundamental, complications.

Clearly this new S-lineage will affect the dynamics of the selection process: by definition it will also compete with one of the two other S-lineages (the one of different S-class from itself). Equally, however, its interaction with the other S-lineage (that of the same S-class as itself) will be quite different, for they must both have the same S-value. The exact nature of the interaction between two such S-lineages of the same S-class (or, more generally, between any two S-lineages of the same S-value) depends on whether procreation is unimodal or multimodal. In the former case, the two S-lineages must still “compete” in the sense that an increase in one must generally be matched by a decrease in the other, but this variation will be purely stochastic—there will be no consistent bias (this is what we *mean* by saying they are of equal S-value). The outcome will be that some one will displace the other in a stochastic manner: this is a pure example of so-called *random drift*. This *may* also be the outcome in the case of multimodal procreation; but in that case there is also an alternative scenario which is that the two S-lineages of the same S-class may completely converge, so that they may *both* survive indefinitely.

At face value, these possibilities make the process of S-lineage selection a little more complicated, but do not seem to have a fundamental significance. From the point of view of the selection dynamics it seems that we could ignore the distinction between S-lineages of the same S-class—because the question of which, if either, of these eventually “wins out” is entirely capricious, and of no particular interest anyway. That is, if S-creation events of this sort

are a common event, we should arguably abandon our analysis in terms of S-lineages, and simply look at the dynamics of the (sub-)populations of the two distinct S-classes *without regard to their parentage*.

This is an innocent looking step, which is implicit in much of modern population (!) genetics—and yet I feel that it critically alters the conceptual structure of the theory. If the step is taken at all (and I admit that, for *practical* purposes in the application of the theory this may well be convenient), it should be only with very careful and explicit acknowledgement of the conceptual implications.

Once one moves to thinking simply in terms of competing sub-populations of differing S-class, one loses the conceptual foundation for the auto-catalytic (Malthusian) growth of these sub-populations, and thus the foundation on which selection is based. Auto-catalytic growth follows from the idea that actors of a given S-class give rise to more actors of that same S-class—but such growth is therefore strictly growth in S-lineages. It may, of course, happen that the sizes of sub-populations also show this kind of auto-catalytic growth—but there is no particular reason why they should; if they do, then this must appear as an extra, essentially ungrounded, property.

Let me put this another way. When we move to thinking in terms of sub-populations, growth in the size of a sub-population may, in fact, be as strongly, or more strongly, related to the size of some *other* sub-population as to itself. The growth, or decay, of a sub-population may, in this case, be nothing at all to do with “selection” in the sense I have been using.

More generally, I will say that, if the qualitative dynamic behaviour of the system *relies* on (rates of) S-creation (i.e. on the generation of actors of one S-class by actors of a different S-class) then the process is of a different sort to the kind I am labelling (Darwinian) selection. It is a central element in what I term selection that the outcome should *not* rely on any details of the process for *generating* “variation” (actors of different S-classes to their parents).

This is actually a very central issue in Darwinian theory. While it is commonly stated that Darwinism is based on “random” variation I think this is a positively misleading formulation. If “random” is interpreted in the sense of the probability calculus, then it implies a predefined and constrained event space. In that kind of scenario it is quite possibly (though not necessarily) the case that the dynamics *do* rely (albeit stochastically) on the probability function of the “variation” process, as much or more than on the properly (in my terms) “selective” pro-

cess.

It seems to me that these issues may be well illustrated in the work of Eigen and Schuster on so-called *quasi-species* (Eigen & Schuster 1979). Briefly, they investigated a model system in which a specific process of stochastic variation was assumed, and they analysed the dynamics on the basis of this. The results rely, not surprisingly, on the details of the stochastic variation. The important point, for my purposes, is that Eigen and Schuster then (in effect?) propose a re-classification of the actors, which “hides” this stochastic variation (yielding what they call quasi-species); and they suggest that, insofar as *selection* is a phenomenon in the system, it is between these quasi-species, and does *not* rely on the details of the stochastic variation, which may be, as it were, factored out.

This impinges on a long-running discussion within evolutionary biology: whether or not the “characteristics” underlying Darwinian selection are required to be “particulate” (or Mendelian).

Neo-Darwinism is, roughly speaking, the synthesis of Darwin’s idea of evolution by natural selection, with Mendel’s theory of “particulate” inheritance. The conventional wisdom is that this augmentation of Darwin’s original theory with Mendelism was essential to the former’s theoretical integrity—that Darwinism simply cannot function in the absence of something more or less like Mendelian inheritance. I do not wish, for the moment, to consider the validity of this view within conventional biological evolution as such; my purpose here is to enquire whether, or in what sense, it can reasonably be claimed that Darwinism cannot function at all, in *any* system, in the absence of (some abstract analog of) “Mendelian” inheritance.

While I am not aware of any extensive analysis of this question, Dawkins (for one) has expressed a clear opinion:

Darwinism post-Fisher is called neo-Darwinism. Its digital nature is not an accidental fact that happens to be true of genetic information technology. Digitalness is probably a necessary precondition for Darwinism itself to work.

Dawkins (1986, p. 115)

Dawkins’ claim for the necessity of Mendelian (“digital”) inheritance is based on contrasting it with so called “blending” inheritance—the idea that an offspring will be, genetically, intermediate between its parents. Blending inheritance leads to the rapid (genetic) homogenisation of a population, and, arguably, evolution by selection would

be impossible—for if variation is absent, or if such variation as may spontaneously appear is diluted before its correlated benefits can be exploited, then it seems that there is no effective basis for selection to operate. Discrete, Mendelian, inheritance, by contrast, does not suffer from such difficulties. Dawkins refers to Fisher’s discussion of this argument;<sup>12</sup> but consider carefully the following comment made by Fisher himself:

The important consequence of the blending is that, if not safeguarded by intense marital correlation, the heritable variance is approximately halved in every generation. To maintain a stationary variance fresh mutations must be available in each generation to supply the half of the variance so lost. If variability persists, as Darwin rightly inferred, causes of new variability must continually be at work.

Fisher (1958, p. 5)

Fisher then argues that the observed mutation rate is not, in fact, sufficient, to maintain the inheritable variability observed in real biological organisms (at least, not under blending inheritance); that if, in fact, blending inheritance were coupled with these observed mutation rates, then it could not support effective evolution; but that, as it happens, blending inheritance is false; and the Mendelian, particulate, theory, when coupled with the given mutation rates, *is* perfectly compatible both with the observed variability, and the effective operation of selective evolution.

I am not convinced that this argument is entirely sound, even as it stands—for example, mutation rates are, to some extent, an outcome of evolution, so a thought experiment (which is, in effect, what is offered by Fisher) that couples blending inheritance with mutation rates produced by evolution based on

---

<sup>12</sup>The original formulation of the argument goes back, of course, to Fleeming Jenkin, though he construed it as an argument against Darwinian evolution in its totality, rather than just against a particular mechanism of inheritance. Given my own professional background, I cannot resist (re-)quoting the following remark made by Darwin’s son, Francis (copied from Hardy 1965, p. 81):

The late Mr. Fleeming Jenkin’s review, on the ‘Origin of Species’, was published in the ‘North British Review’ for June 1867. It is not a little remarkable that the criticisms, which my father, as I believe, felt to be the most valuable ever made on his views, should have come, not from a professional naturalist but a Professor of Engineering.

While I shall still conclude by rejecting Jenkin’s argument, I hope to show that this is a little more difficult, and more significant, than has commonly been supposed.

a different inheritance mechanism is of dubious validity; but I do not propose to pursue this aspect. Even granting that Fisher is entirely correct here, his argument applies only to the biological world as we know it. That is, to turn the passage quoted above back against Fisher (and thus, implicitly, against Dawkins), it could be read as an admission that, if a particular system *does* exhibit a “high” rate of mutation, per generation, then *some* kind of “selective evolution” may become perfectly practical, *even* with a blending mechanism of inheritance! This result is entirely independent of what may, or may not, happen to be the case for terrestrial biology.

More generally, the Fisher/Dawkins argument only goes through in the case of sexual reproduction (multimodal procreation): the contrast between digital and blending inheritance only arises in that case. By contrast, it would seem that some kind of “selective evolution” would be possible based on non-Mendelian (“continuous”) inherited characteristics, even with “low” mutation rates, in the case of *asexual* reproduction. Of course, various arguments might be made regarding the significance of sexual reproduction for the “ultimate” outcome of evolution. In particular, it might conceivably be claimed that a purely asexual form of reproduction would always be limited in some kind of “evolutionary potential” as compared to sexual reproduction—that, other things being equal, a system incorporating sexual reproduction could always “progress” further. But this question is very complex. Dawkins himself has recently commented that the “problem of what sex is good for is still as tantalizing as ever” (Dawkins 1989, p. 274). Clearly, no strong argument for the “necessity” of digital inheritance could be based on so weak a foundation.

In contrast to all this, I want to suggest that Dawkins was essentially right in his original claim—that digital inheritance is necessary for “Darwinian” evolution—but that this is so for quite a different (and more fundamental) reason to that which Dawkins himself cites. It seems to me that any kind of theory of inheritance of “continuous” characteristics (whether sexual or not, blending or not) *must* (implicitly or explicitly) incorporate some specific theory of the *generation* of new variation (whether stochastic or otherwise)—and it is precisely *this* which is problematic. While one can formulate systems of this kind which would exhibit a kind of “selective evolution”, the selection process in question would necessarily rely on, or reflect the characteristics of, the variational process. I think this is unavoidable in the case of continuous characteristics; and *can* still occur even in the case of discrete or

digital characteristics. The most we can say is that, if we want a selection process in which the selection dynamics are effectively decoupled from the variational dynamics then digital inheritance is a necessary (though not, in itself, sufficient) condition for this. And, of course, it is precisely the latter kind of decoupled selection process for which I am choosing to retain the label *Darwinian*.

I do not deny the possibility that processes in which the variational and selective dynamics are coupled may occur. The important point, for my purposes, is that we *can* formulate an evolutionary theory which decouples variation and selection; that we can *only* do this with S-classes based on digital characteristics; that such a theory is quite distinctive and interesting in its own right; and that it therefore deserves to be terminologically distinguished. I feel that this is the kind of theory which Darwin himself hit upon, and that is why I suggest reserving the term *Darwinian* for just this case; but, of course, nothing hinges on what word we use—the distinction is the important thing.

## 6.6 Natural or Artificial?

I have not distinguished between “natural” and “artificial” (S-lineage) selection. The fact is that, within the framework as so far defined there is no distinction. Such a distinction relates solely to the *mechanism* or *process* which maps the S-class of an S-lineage onto its S-value (in given environmental conditions). S-lineage selection requires that some such mechanism exist; but it places no restriction on how it is realised. I shall now add simply that, if this mechanism relies on the intervention of some “intentional agent”, then S-lineage selection is *artificial*, and that otherwise it is *natural*.

*Artificial* selection thus implies that some person (“intentional agent”) must (somehow) examine certain characteristics of the actors, characteristics which are preserved through S-descent, and must then (somehow) control the S-value of their S-lineages, based on these characteristics (for example, by regularly culling all those actors with “undesirable” characteristics, thus increasing the S-mortality of certain S-lineages relative to others). This will cause certain S-lineages to have greater S-value than others, and thus (potentially at least) to displace or eliminate them.

*Natural* selection, on the other hand, means that, although more or less “intentional” agents may be participating in the system in various ways, they are not affecting the relative success in procreation of the actors, or at least are not affecting them in a way

which is systematically related to their S-classes—which is to say that they do not “deliberately” control the S-values of the S-lineages.

The essential question here is whether S-values are being adjusted in accordance with some anticipatory<sup>13</sup> or predictive model of the S-lineage selection process itself, in order to ensure a particular outcome. If so then selection is artificial, otherwise it is natural.

## 6.7 In Summary

In summary, I have equated Darwinian Natural Selection to S-lineage selection—a process whereby one S-lineage deterministically displaces another because it can more efficiently exploit the same resources. It can occur only if several conditions are satisfied:

1. S-class must be predictive of S-fecundity and S-mortality (in the relevant S-environments).
2. In favourable environmental conditions (i.e. in the absence of resource limitations) both S-lineages must be capable of demonstrating exponential growth.
3. The growth of both S-lineages must be limited by the same resources.
4. The S-lineages must be such that they cannot converge; that is, an increase in the S-size of one must necessarily imply a reduction in the S-size of the other.
5. The expected difference in the rates of change of S-size of the two S-lineages must be in a consistent direction, regardless of the exact S-sizes, and large enough to outweigh stochastic variations.
6. S-creation must be relatively infrequent (compared to the time scale required by the selection episode). That is, the selection process should be decoupled from the variation process. This implies, in turn, that the characteristics underlying S-classification, and inherited through each S-lineage, must be discrete or digital.
7. These various conditions must *not* be met through the intervention of an “intentional” agent.

While I have focussed on the phenomenon of one S-lineage selectively eliminating another, this is not

<sup>13</sup>In the sense of Rosen (1985).

to imply that this is the only kind of phenomenon or behaviour which can arise in the kinds of system I have discussed. Thus, it should be clear that the general framework also supports a variety of other (“ecological”) interactions, for example:

- S-lineages may rely on completely independent resources, in which case each can thrive (or go extinct) independently of the other(s).
- One S-lineage may effectively *be* a resource for another—which is to say they are in a predator/prey, or perhaps a parasite/host, relationship.
- Several S-lineages may positively co-operate with each other—where each assists or facilitates the success of the others, which is to say, they are symbiotes.

Nonetheless, it is selection which is of central importance in Darwinian theory, and it is this which I shall continue to emphasize.

In principle, S-lineage selection can continue indefinitely—either due to changes in environmental conditions, or to the birth of new S-primaries. In the latter case, if a new S-lineage is more efficient at utilising some resource, it may selectively displace some prior S-lineage (which, presumably, had been the most efficient up to that time).

This discussion of selection has been informal and qualitative. It can, of course, be given an elaborate mathematical basis—see, for example, (Maynard Smith 1989) or (Eigen & Schuster 1979). This allows analysis, in the simpler cases at least, of the circumstances in which S-lineage selection will occur, how long it will take, and whether it will be partial or total etc. However, these detailed issues are not essential to my purposes here. The point I want to make is that S-lineage selection arises because S-descent preserves S-class (thus ensuring that S-lineages consist of actors which are all members of one S-class), and because S-class *also* characterises certain relevant attributes of the members (specifically, whatever attributes have a significant effect on the fecundity and mortality of the containing lineage), so that the S-lineages can consequently display deterministic behaviours of growth, competition and, ultimately, selection.

## 7 Darwinian Evolution

The notion of Darwinian *evolution* is already largely implicit in my discussion of Darwinian *selection*, but I should like to now make it fully explicit.

Clearly, in any system of actors where S-lineage selection takes place, there will be a change, in time, of the characteristics of various lineages which contain the S-lineages. That is, even though S-creation identifies a discontinuity in an S-lineage, the resulting S-primary will, in general, still be an offspring of some pre-existing actors; it will be a member of some pre-existing lineage, though it also simultaneously founds a new S-lineage within that lineage. Within such a lineage then, we may have multiple S-lineages competing with each other and, over time, we may have multiple episodes of S-lineage selection, so that the actors at a later time are representatives of S-classes which are very different from the S-classes represented at an earlier time. Thus, lineages can be said to *evolve* via a process of S-lineage selection.

*Darwinian* evolution is then a process of evolution by natural selection between S-lineages, but with one further constraint: the S-creation process must be “unjustified”.

Now I have earlier stipulated that the S-lineage selection process should be decoupled from the S-creation process (i.e. from the generation of variation). I am now adding a further, separate and independent, constraint on the S-creation process. By requiring that S-creation be “unjustified” I mean that it must not involve the application of anticipatory or predictive models of the S-value of the resulting S-lineage.

Note carefully that this constraint is not a blanket exclusion of teleology (in the sense of anticipatory behaviour, or *knowledge*). It may well be that individual actors and/or S-lineages incorporate knowledge of their environments; indeed, this possibility is what motivated my examination of Darwinism in the first place. The Darwinian constraint is simply that the course of evolution must not be *directed* by anticipatory models; all other aspects of system behaviour can involve anticipatory models, or knowledge. Thus, the requirement is either that the actors simply do not have anticipatory models for evolution (the normal case), or if they *do* have such models, that these models are not employed in order to *control* evolution.

I shall refer to actors satisfying all the conditions for participating in a process of Darwinian evolution as Darwinian actors, or simply *D-actors*.<sup>14</sup> I shall call a system comprising D-actors embedded in an environment which permits Darwinian evolution to occur a *D-system*. That which evolves by Darwinian means is then, precisely, a lineage of D-actors, or a *D-lineage*.

<sup>14</sup>Gould (1982) has previously introduced the term “Darwinian actor” more or less as I use it here.



## 8 Climb Every Mountain

The evolutionary literature is literally jammed with representations of selective surfaces whose construction is obscure . . . It should give pause to consider that for over fifty years the majority of evolutionary biologists have believed Wright's 1932 diagrams of the adaptive landscape to be among the most heuristically valuable diagrams in all of evolutionary biology, yet to discover that the surface as he conceived it is unintelligible.

Provine (1986, p. 316)

The idea of a "selective surface" is roughly this. Suppose that we represent every possible S-class by a point in a finite dimensional space. Since S-value is associated with S-class, we can associate an S-value with every point in the space. Let us introduce one further dimension to represent this. The result is then a "surface" of S-value. The significance of this surface is that Darwinian evolution can be viewed as a process in which S-lineages are progressively displaced by other S-lineages which are at a "higher" level on this surface; which is to say that Darwinian evolution can be viewed as a form of *hill climbing*.

This view is startling, and perhaps even beautiful, in its simplicity. It is hard to resist the idea that this is a deep and powerful insight into the nature of Darwinian evolution. Unfortunately I want to suggest that, in all but the very simplest of situations, this view is at best incoherent or, at worst, entirely mistaken.

I should emphasize that my representation of the hypothetical fitness surface here is not quite the same as any of the several variations on the idea considered by Provine (1986, pp. 307–317). Furthermore, while I share Provine's general skepticism as to the utility of the fitness surface concept, my objections to it are somewhat different, and perhaps more deeply seated, than his.

A first objection is as follows. If the argument already given for the discrete or digital nature of S-classification is accepted, then it follows that the graph of S-value will be a set of discrete, disconnected, points and not a "surface" (in the sense of a continuum) at all. This is true as far as it goes, and underlies (at least roughly) Provine's criticism of Wright's earliest published version of the surface; but it seems to me that this is really not too serious an objection. The discrete nature of the selective "surface" can be admitted while still retaining the basic image of Darwinian evolution as a hill climbing

process (albeit the steps must be viewed as necessarily discrete).

By contrast, my substantive objection to the fitness surface idea turns on the very notion of S-value, and whether a "graph" of S-value against S-class is a coherent or helpful notion at all.

Recall that S-value was defined as a quantity relating two S-lineages rather than as a property of any single S-lineage in itself. This immediately raises the question of whether it is meaningful to talk of a graph of S-value against S-class in any general sense. My view is that this idea *can* make sense: but only in circumstances that are so circumscribed as to be of very limited interest, at best. Even then, the notion of "hill climbing" as such hardly arises.

In outline, I suggest that we can meaningfully speak of a selective "surface" only in relation to a set of S-classes such that S-lineages of these S-classes would all be pairwise competitive (that is, each such pair meets *all* the requirements for selection, which implies that they are all reliant on the same resources and are all more or less exclusive of each other); furthermore, the resulting S-values must satisfy some kind of approximate, if not exact, transitivity relation (if A can displace B, and B can displace C, it must be the case that A can displace C). If we can identify *all* S-classes which would compete with each other for a particular set of resources (i.e. all S-classes which might potentially compete to occupy a particular environmental *niche*) then it is fair to view Darwinian evolution as a process in which each S-lineage occupying this niche may be progressively displaced by an S-lineage of higher S-value.<sup>15</sup>

Even this is not "hill climbing". Hill climbing as a notion only makes sense in a space where the "steps" satisfy some condition of "locality". If such a locality condition holds, then hill climbing will involve proceeding to the "nearest" peak from the initial point, and staying there (this may or may not be a global peak). But I have been at pains to emphasize that, in my presentation of Darwinian evolution, *no* assumption is made about the mechanism or direction of the variational process. Thus, it will be generally the case that, for any particular representation of the S-classes in question (i.e. any geometry for the space of S-classes to which the selective surface will be added) variations will *not* be, in any sense, "local", and the evolutionary trajectory, although it would consistently transit through S-lineages of higher S-value, would not proceed to

---

<sup>15</sup>All this relies on an assumption of continued environmental constancy, of course: otherwise the S-values themselves would be dynamically changing also.

the “nearest” peak, and might (or might not) cease at any point (regardless of whether that is a “local” peak—the only exception being the global peak, where Darwinian evolution would still necessarily cease).

So, even in very restricted circumstances, the hill climbing view is hardly meaningful. In the general case, it becomes wholly untenable. Firstly, of course, the assumption of environmental “constancy” cannot be upheld over long evolutionary periods. But secondly, and much more seriously, it seems to me that it is quite mistaken to suppose that distinct environmental niches should be equated to distinct peaks in some kind of (suitably defined) selective surface. On my analysis, distinct niches would correspond rather to distinct S-class spaces—with entirely distinct selective “surfaces”. While S-values in these different spaces might be expressed in the same units, they would be basically incommensurable. S-lineages from different such spaces would not compete with each other, and taking the difference in their S-values is a nonsense.

Of course, in any real application of Darwinian theory we cannot expect to see neatly delimited, constant, niches anyway. “Pure” selection (where an S-lineage actually completely displaces one or more competitors) may be the exception rather than the rule. Furthermore, any given D-actor may simultaneously be a member of many S-lineages occupying distinct “niches” and competing with rivals over radically different time scales.<sup>16</sup>

In the general case, niches exist and are defined only relative to the S-lineages exploiting them. In this sense, Darwinian evolution must encompass a dynamics in which the distinct S-class spaces associated with distinct niches will be coupled to each other, and will themselves undergo some kind of ongoing structural transformation (so that the notion of a static S-class “space”, of finite dimensionality, is, at best, only a temporally limited approximation). Popper (1979, Essay 7) has emphasized the process whereby a “new” niche becomes defined as a crucial aspect of Darwinian evolution (though he does not describe it quite in the terms used here). Lewontin (1983) has aptly suggested that, in place of seeing Darwinian evolution as steadily ascending some “selective surface”, we should think rather in terms of a walk on a trampoline; this is an improvement, though I think it may still be a gross oversim-

<sup>16</sup>In this way, it may be that some forms of so-called “species selection” should be seen as identical to conventional “organism selection” (they are both properly called S-lineage selection, with organisms as D-actors) except that the timescale of the selection episodes is much longer for the former.

plification.

## 9 On Adaptation

I have made no mention of “adaptation”, or of the growth of knowledge (in the most general sense), in this abstract formulation of Darwinism. At first sight this seems strange—for it is typically just such phenomena which we wish to study (or even *explain*) by invoking Darwinism.

I suggest that almost the best we can say about this is that D-systems evidently realise a form of unjustified variation and selective retention, in the sense of D.T. Campbell (1960a; 1960b; 1974a);<sup>17</sup> as such, they may *permit* the growth of knowledge (including “inate” knowledge, or “adaptive complexity”) but cannot *guarantee* it.

Thus, knowledge *may* spontaneously grow in a D-system, but it also may *not*; indeed, knowledge previously gained may even be lost. If knowledge ever does grow, it will do so only when the following additional conditions are satisfied:

- S-creation (“blind variation”) must be such that at least some S-lineages can have greater knowledge than any which already exist.
- Knowledge must be well correlated with S-value for at least some S-classes—i.e. some knowledge must be heritable *and* selectively advantageous.

It may seem unnecessary to stipulate, as an explicit constraint or hypothesis, that additional knowledge confer selective advantage—for surely an S-lineage which “knows” more than its competitor(s) will inevitably have higher S-value? This sounds plausible, but in fact it would only hold to the extent that knowledge has no *costs* associated with it; since, in general, the embodiment of knowledge may well impose costs upon S-lineages, it may be that the more knowledgable are *not* favoured by selection.<sup>18</sup>

To put it another way, the only knowledge or adaptation whose acquisition can be explained in Darwinian terms, is knowledge which demonstrably serves the self-interest of the S-lineage possessing it: if there are things that it is better for an S-lineage

<sup>17</sup>At least, I hope this is a correct assessment—since my formulation of the D-system concept was motivated precisely with this result in mind.

<sup>18</sup>It may be argued that the substantive distinction between the growth of knowledge in general and the growth of conscious scientific knowledge resides precisely in the fact that, in the latter case, there is a deliberate or conscious attempt to *ensure* that increases in knowledge confer selective advantage. However, I shall not pursue that argument here.

not to know (since they do it no *net* good in its battle with rival S-lineages) then so be it, and ignorance will in fact be bliss. Even if, by chance, S-creation should give rise to a new S-lineage which is considerably more knowledgable than any of its predecessors, it will be promptly eliminated unless this extra knowledge actually gives it an edge in resource utilisation over relevant competitors.

Thus, it is only if the attribution of “D-system” is coupled with an adequate theory of why, in the particular system, the more-knowledgable S-lineages are also the more successful, and a theory of how more-knowledgable S-lineages can be spontaneously born (without teleology), then this *complete* theory may explain some particular episodes of the growth of knowledge (by Darwinian processes) in this (D-)system.

We might be tempted, at this point, to seek some further augmentation of the concept of D-system, which would *guarantee* the spontaneous growth of knowledge—presumably by guaranteeing that the additional conditions stipulated above will be satisfied. However, it seems to me that it would be a mistake to embark on such a project. For the conditions, as described, are only meaningful in the context of some *particular* knowledge; given some such specification we might well be able to establish that *that* particular knowledge could (or could not) spontaneously emerge, by strictly Darwinian processes, in some particular system; but we could never examine a system and say simply that “knowledge” (of an unspecified nature) will grow in it. The best that we might hope for is the negative formulation that a given system is such that the knowledge embodied within it (if any) definitely *cannot* grow (for example by demonstrating that no more-knowledgable S-lineages are accessible via S-creation, or that no accessible increase in knowledge is selectively advantageous). I doubt however that even this negative conclusion could be confidently asserted for any but the simplest systems.

More abstractly, I suggest that the search for a *general* theory which could “predict” the growth of knowledge (whether by Darwinian means or otherwise) would be tantamount to the search for a *logic of induction*, and would be equally futile. The strictly limited scope of the purely Darwinian theory, and its close relationship to the problem of induction, is discussed by Popper (1976, Chapter 37) (albeit in the purely biological interpretation of Darwinism). He suggests that Darwinism *per se* may be best viewed as a *metaphysical* rather than a *scientific* theory. This should not be taken as pejorative: Popper himself describes the Darwinian theory

as having been “invaluable”; and this is because it provides a very general *schema* according to which particular, properly scientific, theories can be developed and, perhaps, tested, in specific cases. It is a schema for that most difficult class of problems: those involving the growth of knowledge. It is a virtually unique tool in this context because it offers the *possibility* of viable, yet *non-theistic*, theories. As Popper says, “. . . theism was worse than an open admission of failure, for it created the impression that an ultimate explanation had been reached”.

I should add that Popper has, since his original analysis, significantly qualified his views on the status of Darwinism—even going so far as to describe his more recent work in this area as a “recantation” (Popper 1978). I take the view that this modification of Popper’s views is concerned, not with the abstract ontology of Darwinism, as discussed in this essay, but rather with the specifically biological application of Darwinism. In any case, it seems to me that Popper’s original analysis has commonly been misunderstood, and that his “recantation” is actually more in the nature of a clarification; as such, I claim that the interpretation I have given above is consistent both with Popper’s original analysis, and with his later, clarified, version. For example, Popper has reiterated that he still believes that Darwinism (“Natural Selection”) works as a research program. Of course, in the final analysis, the arguments given in this essay must stand (or fall) on their own merit, and not on the implied authority of Karl Popper (or anyone else).

Note that it follows from the discussion that there is nothing in the strictly Darwinian framework which explains (much less guarantees or predicts) any particular *progress* (in the sense of the growth of knowledge at least) in Darwinian evolution. At its simplest, this can be recast as being a consequence of the fact that S-creation and S-lineage selection are devoid of anticipatory or teleological influences or controls: there is thus no *a priori* direction in Darwinian evolution, though we may be able to recognise “progress” after the event. In particular, while knowledge may grow for an indefinite period, there is never any guarantee that growth will continue, or even that the knowledge already achieved will not be lost; as Popper notes “. . . there is more than one way in which all life on Earth might be destroyed”, (Popper 1976, p. 168).

Conversely, we must not suppose that, because no system can be *guaranteed* to exhibit, or continue to exhibit, the growth of knowledge, then all D-systems are “equal”; they are not. While it is futile to seek general conditions which would be *sufficient* for the

growth of knowledge, we may well be able to identify further conditions which are necessary, or even which are only beneficial in some cases (provided they are not positively detrimental in others).

So: Darwinism is no “ultimate explanation” but it offers the possibility of particular, local, and limited explanations. This is not a criticism of Darwinism: for if the analysis given here is valid then “ultimate explanation” is not possible. Darwinism may be “wrong” (in any particular case of the growth of knowledge), and there may yet be preferable theories (or research programs); but this preference will not be by virtue of another theory being logically “stronger”, for Darwinism is already as strong as its problems permit.

## 10 M.B. Williams: Darwinian Axiomatics

Mary B. Williams has attempted a detailed, formal, *axiomatisation* of Darwinian theory (Williams 1970). Although I originally developed my own framework, as presented above, largely in ignorance of Williams’ work, it transpires that there are substantial similarities in our approaches (convergent evolution perhaps?)—and some significant differences also.

Firstly, there are superficial differences of terminology and notation, some of which have already been noted. Thus, for example, Williams uses *Darwinian biocosm* where I use *D-system*; or, again, Williams has *biological entity* where I have *D-actor* etc. At this level our terms are largely interchangeable.

In addition to a special technical vocabulary, Williams introduces a formal mathematical notation, which is quite austere and concise, because she wishes to present formal mathematical *theorems* derived from her axiomatisation of Darwinism. My treatment, on the other hand, is essentially *informal*, and more concerned with identifying an appropriate ontology—a set of more or less revealing concepts which underlie Darwinism—rather than the formal establishment of theorems. Thus I have not introduced a separate concise, mathematical, notation.

I now turn to more substantive issues.

Williams and I have had similar objectives in that we have both explicitly attempted to formulate a more or less *abstract* version of Darwinism—for the reason that this does not prejudge which (if any) aspects of reality might be well modelled by it. However, Williams’ motivation is still primarily

the application of Darwinism in a biological context. In particular, she wishes to argue for the possibility that the abstract Darwinian theory might have more than one separate, concurrent, and mutually compatible, interpretation even *within* the biological world. I do not dispute this possibility, but wish also to emphasize the possibility of applying Darwinism outside areas conventionally thought of as biological—such as the growth of knowledge, especially in artificial systems. This incidentally explains why Williams continues to use some explicitly biological allusions (*biocosm*, *biological entity* etc.) whereas I have generally tried to avoid this.

Williams emphasizes, as I do, that it is *lineages* (she calls them *clans*) and not individuals which are selected between. In Williams’ presentation, my notion of S-descent and S-lineage corresponds to the existence of subclans within subclans, or *subclands*, as Williams puts it, which are distinguished from each other by the fact that they selectively compete with, and ultimately eliminate each other. This requires the postulation of some attribute(s) which characterise a subcland (i.e. only those offspring are included in a subcland which inherit the attribute(s) characteristic of the subcland) and which are such that they jointly determine the outcome of selection between distinct subclands. Williams conflates all the relevant attributes into a single parameter which she terms the “fitness”. As already discussed, I have chosen to avoid the word “fitness” since it is potentially confusing, and have chosen instead to introduce the distinct term *S-value*.

It is worth emphasizing that Williams’ fitness is directly axiomatised (*Axiom D3*) as a property of individual biological entities (D-actors), whereas my S-value is specifically introduced as a characteristic of an S-lineage (albeit implied by the S-class of its D-actors). I shall argue below (section 11) that although these two approaches are, to some extent, interchangeable, it is conceptually much simpler to attach the idea of fitness (or S-value, or whatever) to S-lineages rather than to D-actors. This is implicit even in Williams’ treatment for, having introduced fitness as a property of biological entities (D-actors) she immediately (definitions *D3* and *D4*) introduces a notion of “average” fitness which can be attached to subclands (S-lineages) and it is this which is used exclusively in the subsequent development. Even when attached to S-lineages, there are further detailed differences between Williams’ concept of fitness and my S-value; however, for the rest of this discussion I shall ignore them, and consider my S-value to be more or less synonymous with Williams’ fitness.

On another issue, Williams implies that selection can only take place under some assumption of *environmental constancy*—for D-actors may be classified alike, and this may be inheritable, but it can be predictive of the outcome of selection only if the offspring are faced with relevantly similar environmental conditions. However, Williams does not actually stipulate environmental “constancy” as such; instead she allows that environmental changes may occur, but stipulates that relative “fitness” must be more or less consistently maintained by some subclands (howsoever they may be identified) for long enough that selection can work itself through: this is the substance of her complex and subtle *Axiom D5*. I suggest that although this appears to allow for arbitrary environmental variation, it actually amounts to a requirement that the environment be constant *in some relevant (exploitable) respects*—at least while selection is actually going on. In any case, whatever about the fine print, the essential point which I accept is that my notion of S-classification can only make sense *relative* to some constraints on the environment of the D-actors.

Williams emphasizes, as do I, that the Darwinian theory “does not depend on any particular theory of inheritance”, such as Mendelism. The *only* substantive constraint on the nature of inheritance is it must be able to preserve distinctions in S-value through indefinite numbers of generations. This is the essence of my requirement that S-descent *preserve* S-classification (coupled with environmental constancy). I have argued that a requirement for inheritance of discrete or digital characteristics flows from this. However, it must be carefully noted that this still does not automatically constrain the detailed implementation of inheritance—for example in the form that it must involve the use of some localisable “hereditary material”.

My final comment on Williams’ treatment is that, in her informal discussion, she occasionally seems to claim more for it than is actually present. In particular, in several places, she informally implies that her axiomatisation *predicts an indefinite and progressive improvement in adaptation*. Yet, according to my discussion previously, if “adaptation” is thought of in the substantive sense of “knowledge” (inate to the D-actors and/or S-lineages), then Darwinism in itself *cannot* predict its progressive improvement—indeed I have argued that that may be too great a demand for *any* theory to bear.

The resolution of this apparent paradox consists of a minor and a major point.

The minor point is covered by Williams herself. Although her theory *does* correctly predict an in-

definite progression in S-value, this is so *only* because the axiomatisation guarantees that S-lineages of higher S-value will *always* be present, in the process of being selected (*Axiom D5*). This is effectively a somewhat arbitrary, and unanalysed, constraint on S-creation. It in no way rules out the possibility that, in any real system, such S-lineages of higher S-value may cease to appear at any stage, and S-value would then stagnate, at best. In such a case, Williams would simply say that the system has ceased to be *Darwinian* (i.e. has ceased to satisfy her axioms). There is no sense in which Williams claims to have provided a general specification of conditions which would guarantee that a system would *always* (or even ever) bring forth S-lineages of higher S-value; her analysis is limited to saying what will *subsequently* happen if this condition is, in fact, satisfied (namely that the S-lineages of higher S-value will fix, and S-value will therefore progressively and indefinitely increase).

The second and more substantive point concerns what Williams means by “adaptation”. Williams does not give any explicit explanation of her usage. Moreover, it does not appear *at all* in the formal terms of her axiomatisation (either as a primitive or in a definition)—and, of course, it cannot then appear in any of the formal theorems. I suggest that Williams (consciously or otherwise) uses the notion of adaptation as being synonymous with (her technical version of) “fitness”, which, in turn, is comparable to my S-value. In this case, it is clear that her conclusion does not bear upon, and therefore cannot conflict with, my analysis of the growth of knowledge.

If, on the other hand, Williams actually intends adaptation to imply something analogous to my meaning of “knowledge”, then I claim simply that the statements she makes in terms of adaptation (for example, that “a continuing sequence of adaptive modifications is ubiquitous in a Darwinian subclan”) do not (indeed could not) follow from her axioms; and that, if this was her intended claim, it could only mean that she overlooked the fact that there is no *inherent* reason why adaptation (still in the sense of inate knowledge) should be correlated with fitness or S-value.

## 11 The Unit of Selection

The question I propose to address in this section is: what entities, or kinds of entities, are the “units of selection”?

This issue has received extended consideration, as summarised by, for example, Hull (1981). Not unnaturally, this discussion has been carried out mainly in the particular context of biological evolution, but Hull remarks that the issue is so fundamental that it deserves to be called ‘metaphysical’. Similarly, Dawkins has suggested that the question is not simply one of empirically deciding which real entities function as units of selection, but rather is a “dispute about about what we ought to *mean* when we talk about a unit of selection” (Dawkins 1982b).

Needless to say, given the intense scholarly attention which has been devoted to this issue, I do not claim to have a definitive resolution to offer. However, I shall at least try to offer a clear target for further criticism.

In particular, I embrace Dawkins’ view that the substantive question here is to elucidate “the nature of the entity for whose benefit adaptations may be said to exist” (Dawkins 1982a, p. 81).<sup>19</sup> This, whatever it is, is the entity which I shall identify as the “unit of selection”.

In my formulation of Darwinian evolution there are really only two candidates for the rôle of the unit of selection: the D-actor or the S-lineage. I shall suggest that, in certain, simple, cases, there may be very little difference between the two; and that this explains why the two possibilities are commonly confused. However, I shall conclude that, in the *general* case, the two possibilities are very different; and that the *correct* candidate is, quite unambiguously, the S-lineage.

On my analysis, if “adaptation” evolves at all as a result of Darwinian processes (and, to reiterate yet again, there is never any guarantee that it will) then it does so in the form of a change (or accumulation of changes) which allows a new S-lineage to selectively displace an old one because it has higher S-value. That is, adaptations will be such as to favour an S-lineage relative to its competitors. This immediately suggests the view that adaptations can be viewed as Darwinian *only* to the extent that they are beneficial to the S-lineage exhibiting them (at least for the duration of the selection episode): i.e. that the S-lineage is the “unit of selection”. However, this needs to be considered in somewhat more depth.

---

<sup>19</sup>There is, of course, an implication here that we are talking exclusively about “Darwinian adaptations”: that is, “adaptations” (whatever they may be) brought about by Darwinian evolution.

In the simplest case, S-value can be improved by effects which are uniform and consistent for all members of the S-lineage. This kind of effect is particularly easy to envisage, given that all D-actors in a given S-lineage are already known to be “similar”, in the sense of having the same S-class.

Ultimately, this effect would be manifested in one of just two possible ways: if all the D-actors uniformly<sup>20</sup> live longer than the D-actors of the competing S-lineage(s), then the S-mortality of the S-lineage will be less than its competitors; or if the D-actors all uniformly have more S-offspring (per unit time) then the S-fecundity of the S-lineage will be greater than its competitors. This way of thinking leads to associating S-value with individual D-actors, at least as a form of shorthand. This is always legitimate in the formal sense that, as already noted, all D-actors in an S-lineage share a single S-class—so that any attribute of an S-lineage (such as S-value) can be imputed to an S-class, and thus, implicitly, to single D-actors, considered as “exemplars” of their S-class.

To put it a slightly different way, from the point of view of any single D-actor, it “knows” for sure that it is a member of its own S-lineage(s). Thus, it can surely favour its own S-lineage(s) by having as many offspring as it can (involving some kind of trade-off between living as long as it can while procreating as fast as it can). Again, this viewpoint estimates or measures S-value by reference to the activities or attributes of any single “typical” member of the S-class.

The problem with this point of view is that there may be no such thing as a “typical” D-actor of a given S-class. It is perfectly possible for D-actors of the same S-class to be very different in a wide range of characteristics or attributes; or, for that matter, for any single D-actor to engage in a wide variety of different behaviours in different times and circumstances; we require only that these differences not be such as to establish two distinct, competing, S-lineages with distinct S-values.

Indeed, the S-value of an S-lineage may actually critically rely on the distribution of variations in certain characteristics among the D-actors making it up, so that no single D-actor, *qua* D-actor, would allow a determination of S-value—except in the contrived sense that we should be able to determine its S-class, and thus determine the “typical” distribution of variations within an S-lineage made up of D-actors from that S-class, and thus, finally, deter-

---

<sup>20</sup>“Uniformly” implying that each D-actor is associated with an independent random sample from a single probability distribution, as previously discussed in section 6.1 above.

mine the S-value of the S-lineage (relative to specified competitors).

We can envisage a variety of ways in which variation in the characteristics of D-actors (within an S-class) could be beneficial to an S-lineage. For example:

1. If there is a diversity of resources or habitats, whose availability varies, then an S-lineage which keeps its options open—distributes its D-actors over these resources or habitats—may well do better than competitors which lack this flexibility (whose D-actors are more uniform). This can work without any assumption of interaction or mutual recognition between the D-actors making up a single S-lineage.
2. Alternatively, if there *is* any mechanism whereby a D-actor can identify other D-actors with which it shares an S-lineage (with some degree of probability) then there may be other opportunities for *cooperation* between D-actors, perhaps involving specialisation or division of labour.

In general, both these possibilities could equally be based on dedicating certain D-actors to certain tasks on a lifetime basis, or on the use of a structured life cycle with different specialisations on the part of a single D-actor at different stages in its cycle; and, of course, the two possibilities could both be exploited to greater or lesser extents in a single S-lineage.

The notion of a D-actor being so constituted (in structure or behaviour) that it functions for the benefit of other D-actors, at a cost to itself (especially as measured in terms of the number of offspring it produces), is referred to as *altruism*, and has historically been viewed as somewhat problematic for Darwinian theory. However, we can now see that this derives from a blinkered view of selection: the view that supposes that a D-actor can benefit its S-lineage(s) (improve its S-value) *only* through direct “benefit” to itself. In particular, the second scenario mentioned above, where an S-lineage can benefit from mutual co-operation between its members, clearly allows the possibility of altruism (on the part of D-actors to other D-actors in a single S-lineage) being *favoured* by a Darwinian process.

I have suggested that altruism could be favoured where there is a mechanism for D-actors to recognise other D-actors sharing a single S-lineage. This is a necessary condition. By definition, altruism involves costs to the D-actors engaging in it. This can conceivably provide a net benefit to the S-lineage *only* if most of the benefit of the altruism is retained within

the S-lineage. To put it another way, altruism is in permanent danger of being subverted or exploited. If one S-lineage contains D-actors which are unconditionally altruistic, and another is otherwise similar (and thus competing) but its D-actors succeed in restricting their altruism (even with limited effectiveness) to members of their own S-lineage, then the latter S-lineage clearly gets more benefits from altruism than its competitor, and will be favoured by selection.

Note that this argument would not generally go though if we considered a putative S-lineage which abandoned altruism entirely. Such an S-lineage could certainly *initially* grow at the expense of the original S-lineage of altruistic D-actors. But recall our basic hypothesis that the altruism of the D-actors was beneficial to the S-lineage. The new S-lineage, of entirely selfish D-actors, would, as it becomes more numerous, start losing this benefit. There is no general answer as to the final outcome of such a process, but it certainly does not follow that the original S-lineage would necessarily be eliminated.<sup>21</sup> It would be equally possible that another different S-lineage would arise which *restricted* its altruism: this would then generally be capable of selectively eliminating *both* of the others.

So altruism, if it arises, is expected to rely, to a greater or lesser extent, on mutual recognition between the D-actors of a single S-lineage. There are two basic mechanisms which suggest themselves for such recognition:

1. One’s own “close” relatives (parents, sibling, offspring) are likely to be members (and, to a lesser extent, their close relatives in turn).
2. One may attempt to “recognise” other members of some S-class.

The former corresponds to what is conventionally referred to as “kin selection” in biological Darwinism, (Dawkins 1979). The latter is not generally explicitly emphasized in biology, though Dawkins has considered a variety of mechanisms which might come under this general heading (Dawkins 1982a, Chapter 8). Notwithstanding this apparent lack of emphasis among biologists, I speculate that benefits of altruism (in the broadest possible sense) mediated by recognition of S-class and/or recognition of close S-kin may be quite significant for Darwinian evolution generally. I say this simply on the basis

---

<sup>21</sup>That is, it is not generally the case that such a process would meet all the specified requirements for selection to go through; in particular, it may not meet the requirement of a *consistent* selective bias (independent of S-size).

that recognition of this sort is clearly, in the general case, a difficult problem that calls for complex sensory processing, i.e. the utilisation of relatively complex anticipatory models of the world. In terms of the argument of section 9 above, benefits (for S-lineages), deriving from more and more effective S-class recognition and discrimination, seem to provide one plausible general basis for a sustained correlation between S-value and knowledge, and thus for an extended period of growth of knowledge through evolutionary time. However, intriguing as this possibility may be, it is of limited relevance to my purposes here, and I shall not pursue it further.

To return to the central idea of this section, which is the notion of the S-lineage as the unit of selection, the significance of this point can be put at its most stark by asking what is the “adaptive” value to a D-actor of procreation itself? Basically, there is none, and yet everything else that might be an adaptation of a D-actor is typically identified in terms of its contribution to procreation. Conversely, there is no problem when we think of the structure and behaviour of D-actors, insofar as it has been adapted by Darwinian evolution, as being consistently for the benefit of the S-lineage. Which is simply to say that S-lineages which have come about through Darwinian evolution, are expected to be characterised by naked and exclusive *selfishness*; altruism on the part of an S-lineage (as opposed to the part of its component D-actors) could only be classified as pathological in a Darwinian context.

The essential point is that while the adaptations of an S-lineage *may* be clearly manifested in “adaptations” of the D-actors (or, more generally, of the S-classes, identified in some particular manner), this is an incidental effect, which may or may not occur, rather than being an intrinsic feature of Darwinian evolution in general. The general and reliable feature is the selfishness of the S-lineages themselves. The failure to realise this is manifested in the attempts to introduce entirely counter-intuitive modifications to the notion of “fitness” of a D-actor when it is discovered (as, for example, in the case of kin selected altruism) that, without such a correction, entities of lower “fitness” are apparently favoured by selection. This is the general idea of so-called “inclusive fitness”. Dawkins has rightly pointed out that that is an inspired, but ultimately misguided, stratagem (Dawkins 1978): the supposed difficulty simply does not arise when one realises that the D-actors are not the appropriate entities to ground the analysis in the first place—the S-lineages are.

So I consider that the unit of selection is the S-lineage; and that the implication of this is that a

given phenomenon can be considered as an outcome of Darwinian evolution if and only if it is consistent with an unconditional selfishness on the part of the S-lineage possessing or exhibiting it. This is the doctrine of *The Selfish S-lineage*.

## 12 Conclusion

The central concept which I have introduced is that of the S-lineage as the unit of (Darwinian) selection. The S-lineage is a peculiar hybrid of individual and class—perhaps it should best be regarded as an ontological category of its very own. It is a bastardized entity and not at all pretty; but it seems to me that this somewhat uneasy alliance of the permanent and the ephemeral is an intrinsic aspect of Darwinian theorizing, and might as well be pulled out into the open so that we may examine it properly. At worst, if I am thoroughly mistaken in this view, I hope that I have at least presented my ideas sufficiently clearly that they can now be criticised and corrected.

As indicated in the introduction, this essay has involved the introduction of a range of new, and occasionally cumbersome, terminology. My claim is that this is justified by the fact that it will provide an effective framework for the application of Darwinian concepts outside the specific field of biology. Admittedly, this claim can be properly tested only in subsequent work.

This is a internal Technical Report, and the freedom of that format has perhaps been abused in the protracted nature of some of the discussion. Nonetheless, I would greatly appreciate comments and criticism of any sort.

## Acknowledgements

This essay is part of an ongoing attack on the problem of realising the spontaneous growth of Artificial Knowledge by Darwinian (or any other!) means. In this pursuit, I have benefited greatly from discussions with colleagues, particularly Noel Murphy in DCU, and John Kelly of University College Dublin. I am also indebted to the School of Electronic Engineering in DCU (particularly through the agency of its Head, Charles McCorkell) for continuing encouragement, not to mention material support. All errors remain, of course, my own responsibility.



## References

- Bendall, D. S. (ed). 1983. *Evolution from Molecules to Men*. Cambridge: Cambridge University Press.
- Burrow, J. W. (ed). 1968. *Charles Darwin: The Origin of Species*. London: Penguin Books. This is a reprinting of (Darwin 1859), together with some ancillary materials, edited with an introduction by Burrow.
- Campbell, Donald T. 1960a. Blind Variation and Selective Retention in Creative Thought as in Other Knowledge Processes. *Psychological Review*, **67**(6), 380–400.
- Campbell, Donald T. 1960b. Blind Variation and Selective Survival as a General Strategy in Knowledge-Processes. *Pages 205–231 of: Yovits, Marshall C., & Cameron, Scott (eds), Self-Organizing Systems*. Oxford: Pergamon Press. Proceedings of an interdisciplinary conference, 5 and 6 May, 1959.
- Campbell, Donald T. 1974a. Evolutionary Epistemology. *Pages 413–463 (Book I) of: (Schilpp 1974)*.
- Campbell, Donald T. 1974b. Unjustified Variation and Selective Retention in Scientific Discovery. *Chap. 9, pages 139–161 of: Ayala, Francisco Jose, & Dobzhansky, Theodosius (eds), Studies in the Philosophy of Biology*. London: The Macmillan Press Ltd.
- Darwin, Charles. 1859. *The Origin of Species by Means of Natural Selection*. First edn. London: John Murray. Reprinted in (Burrow 1968).
- Dawkins, Richard. 1976. *The Selfish Gene*. Oxford: Oxford University Press. See also Dawkins (1989).
- Dawkins, Richard. 1978. Replicator Selection and the Extended Phenotype. *Zeitschrift für Tierpsychologie*, **47**, 61–76.
- Dawkins, Richard. 1979. Twelve Misunderstandings of Kin Selection. *Zeitschrift für Tierpsychologie*, **51**, 184–200.
- Dawkins, Richard. 1982a. *The Extended Phenotype*. Oxford: W.H. Freeman and Company.
- Dawkins, Richard. 1982b. Replicators and Vehicles. *Chap. 3, pages 45–64 of: King's College Sociobiology Group (ed), Current Problems in Sociobiology*. Cambridge: Cambridge University Press.
- Dawkins, Richard. 1986. *The Blind Watchmaker*. London: Penguin Books.
- Dawkins, Richard. 1989. *The Selfish Gene*. New edn. Oxford: Oxford University Press. This is a significantly revised edition of (Dawkins 1976), with endnotes and two new chapters added.
- Eigen, Manfred, & Schuster, Peter. 1979. *The Hypercycle: A Principle of Natural Self-Organization*. Berlin: Springer-Verlag.
- Fisher, Ronald A. 1958. *The Genetical Theory of Natural Selection*. Second revised edn. New York: Dover Publications, Inc. This is a revised and enlarged edition of the work originally published in 1930 by the Oxford University Press.
- Gould, Stephen Jay. 1982. Darwinism and the Expansion of Evolutionary Theory. *Science*, **216**(23 April), 380–387.
- Hardy, Alister. 1965. *The Living Stream*. London: Collins.
- Hull, David L. 1980. Individuality and Selection. *Annual Review of Ecology and Systematics*, **11**, 311–332.
- Hull, David L. 1981. Units of Evolution: A Metaphysical Essay. *Chap. 1, pages 23–44 of: Jensen, U. J., & Harré, R. (eds), The Philosophy of Evolution*. The Harvester Press.
- Hull, David L. 1983. Darwin and the Nature of Science. *Chap. 4, pages 63–80 of: (Bendall 1983)*.
- Lewontin, R. C. 1970. The Units of Selection. *Annual Review of Ecology and Systematics*, **1**, 1–18.
- Lewontin, R. C. 1983. Gene, Organism and Environment. *Chap. 14, pages 273–285 of: (Bendall 1983)*.
- Maynard Smith, John. 1989. *Evolutionary Genetics*. Oxford: Oxford University Press.
- McMullin, Barry. Forthcoming. *Essays on Darwinism. 3: Genic and Organismic Selection*. Technical Report bmcm9203. School of Electronic Engineering, Dublin City University, Dublin 9, IRELAND.
- Popper, Karl R. 1976. *Unended Quest*. Glasgow: Fontana/William Collins Sons & Co. Ltd. First published as 'Autobiography of Karl Popper' in (Schilpp 1974, Part One, pp. 2–181, Book I).

- Popper, Karl R. 1978. Natural Selection and the Emergence of Mind. *Dialectica*, **32**(3–4), 339–355.
- Popper, Karl R. 1979. *Objective Knowledge: An Evolutionary Approach*. Oxford: Oxford University Press. First published 1972; this edition reprinted with corrections and a new appendix.
- Provine, William B. 1986. *Sewall Wright and Evolutionary Biology*. Series: Science and its Conceptual Foundations. Chicago: The University of Chicago Press. Series Editor: David L. Hull.
- Rosen, Robert. 1985. *Anticipatory Systems*. IFSR International Series on Systems Science and Engineering, vol. 1. Oxford: Pergamon Press. Editor-in-Chief: George J. Klir.
- Schilpp, Paul Arthur (ed). 1974. *The Philosophy of Karl Popper*. The Library of Living Philosophers, vol. XIV. Illinois: Open Court.
- Varela, Francisco J. 1979. *Principles of Biological Autonomy*. New York: North-Holland.
- Williams, Mary B. 1970. Deducing the Consequences of Evolution: A Mathematical Model. *Journal of Theoretical Biology*, **29**, 343–385.
- Wright, Sewall. 1932. The Roles of Mutation, Inbreeding, Crossbreeding and Selection in Evolution. *Pages 356–366 of: Proceedings of the Sixth International Congress of Genetics*, vol. 1. (As cited by Provine 1986).