

Replicators Don't!*

Barry McMullin

School of Electronic Engineering
Dublin City University
Dublin 9
Ireland

Phone: +353 1 704 5432
Fax: +353 1 704 5508
E-mail: McMullin@EEng.DCU.IE
WWW: <http://www.eeng.dcu.ie/~mcmullin/home.html>

Abstract: Replicators don't. Replicate, that is. This is the shocking conclusion to which I have been forced by my attempt to figure out what *precisely* Richard Dawkins means by the term “replicator”. Actually, it seems that Dawkins uses the term in at least two fundamentally different ways; but according to Dawkins' own specification of the problem which the “replicator” concept was intended to solve (namely, what entities can qualify as things that evolutionary adaptations are “good for”) then “replicators” turn out to be a special form of *lineage* (what I shall term a *similarity lineage*); and these, in turn, do not actually “replicate” (in Dawkins' sense of the term) at all! Does this matter to the research programme of Artificial Life? Well yes, I believe it does. Dawkins has explicitly argued that there are principled reasons why Darwinian evolution, *in any medium whatsoever*, must rely on the participation of “replicators”. Within limits I am inclined to agree. But it follows that, if we wish to realize artificial Darwinism, we had better be clear what a replicator actually is—and all the more so if it turns out that it doesn't...

1 Introduction

The notion that “replicators” play a uniquely distinguished role in biological evolution has been championed by Richard Dawkins (1976; 1989). Furthermore, Dawkins has argued that this idea can be generalised in a way which makes it applicable to any properly Darwinian evolutionary process, at least if that process gives rise to a growth in adaptive complexity (Dawkins 1983). It is evident, therefore, that if Dawkins' analysis is correct, it has profound implications for any attempt to realise a growth of adaptive complexity in *artificial* systems by Darwinian means.

This paper is concerned with trying to clarify just what, exactly, a Dawkinsian “replicator” might be. This has implications both for the specific field of Artificial Life, but also for the general debate in evolutionary biology about “units of selection”. I provide a reformulation which, I claim, captures the valid core of Dawkins' insight, while, at the same time, avoiding certain confusions and misconceptions which might otherwise be read into his views.

2 Replicating Confusion

There has been considerable ambiguity, if not downright confusion, in the literature of evolutionary biology regarding specific technical usage of the terms *replication* and *replicator*. As

*Derived from material first presented in a series of three rather indigestible essays, previously available only in the form of internal technical reports (McMullin 1992a; 1992b; 1992c).

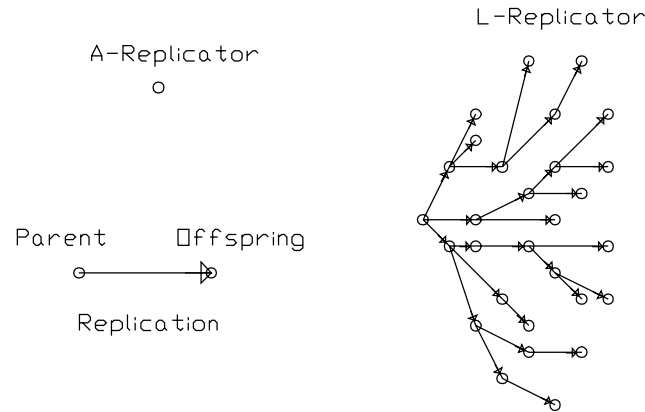


Figure 1: A-replicator vs. L-replicator

far as I am aware, the abstract, technical, idea of a *replicator* was first introduced by Dawkins (1976; 1978a). Hull subsequently elaborated the idea (Hull 1980; 1981), and Dawkins has since extended his own analysis somewhat further (Dawkins 1982a).

I shall argue that “replicator” has sometimes been used to refer to *actors*—individuals which are capable of replication¹—and sometimes to refer to *lineages* formed as a *result* of such replication (Fig. 1). I suggest that the distinction between these two kinds of entities is actually very significant—and that it is only the latter (lineage) interpretation which can do the job Dawkins wants done.

*Where necessary in the following, I shall explicitly distinguish references to actors with an **A-** prefix, and references to lineages with an **L-** prefix.*

The ambiguity of usage can be illustrated by considering the concept of replicator *longevity*. In *The Selfish Gene*, Dawkins first defines longevity as relating to the lifetime of an individual replicator, i.e. an actor or A-replicator:

Certain molecules [supposed primordial replicators], once formed, would be less likely than others to break up again. These types would become relatively numerous in the soup, not only as a direct logical consequence of their ‘longevity’, but also because they would have a long time available for making copies of themselves. Replicators of high longevity would therefore tend to become more numerous and, other things being equal, there would have been an ‘evolutionary trend’ towards greater longevity in the population of molecules.

Dawkins (1976, p. 18)

He goes on to argue that there would be an overall trend toward the evolution of “varieties” (?) of replicator with high “longevity/fecundity/copying-fidelity” (Dawkins 1976, p. 19). Thus, by “longevity” he must evidently mean the lifetime of individual actors, or A-longevity.

Somewhat later in the same source, Dawkins specifies that “Copying fidelity is another way of saying longevity-in-the-form-of-copies *and I shall abbreviate this simply to longevity*” (Dawkins 1976, p. 30, emphasis added). Now *this* version of longevity evidently refers to a replicator viewed as a lineage (“in-the-form-of-copies”); so this is L-longevity.

So far, any confusion is latent: as long as we remember that Dawkins is using “longevity” in two quite different ways, and judge his meaning from the context, it should not cause too much trouble. In particular, we might reasonably suppose that the slogan

¹See (McMullin 1992a, Section 3); *cf.* “Darwinian actor” (Gould 1982).

“longevity/fecundity/fidelity” will always refer to actors, not to lineages—i.e. the “longevity” in question will be A-longevity rather than L-longevity. I say this for two distinct reasons. Firstly, A-longevity is the sense of longevity with which Dawkins first introduced the slogan. But secondly, and more significantly, Dawkins claims that L-longevity is effectively equivalent to copying *fidelity* (a dubious equation in any case, but let it stand). It follows that, if the longevity in the slogan were interpreted as L-longevity, the slogan would become synonymous with “fidelity/fecundity/fidelity”—which is at least redundant and confusing, if not actually incoherent.

Unfortunately, however, Dawkins did indeed subsequently use the slogan in precisely this confusing way:

The qualities of a good replicator may be summed up in a slogan reminiscent of the French Revolution: Longevity, Fecundity, Fidelity [Dawkins 1976; 1978b]. Genes are capable of prodigious feats of fecundity and fidelity. *In the form of copies of itself, a single gene may persist for a hundred million individual lifetimes.*

Dawkins (1978a, p. 68, emphasis added)

So we have the slogan, which I have just argued *must* imply A-replicator, followed immediately by an elaboration that obviously implies L-replicator. Given Dawkins’ own confusion here, it is hardly surprising that Hull then compounded the error further:

According to Dawkins [Dawkins 1978a, p. 68], the qualities of a good replicator may be summed up in a slogan reminiscent of the French Revolution: Longevity, Fecundity, Fidelity. As striking as this slogan is, it can easily be misunderstood. The fidelity which Dawkins is talking about is copying-fidelity, and the relevant longevity is longevity-in-the-form-of-copies [Dawkins 1976, p. 19, p. 30].

Hull (1981, p. 31)

Hull’s last citation here refers to the two locations in *The Selfish Gene* (1976 edition), which I have already identified above, where “longevity” was defined—but he omits to mention that these are two different, *and incompatible*, definitions!

Hull is certainly correct that Dawkins’ slogan may be easily “misunderstood”. On my view, both he (and Dawkins himself) have suffered from just such misunderstanding. The interpretation Hull gives here is not the “correct” one—i.e. that which accompanied the original formulation of the slogan in (Dawkins 1976, p. 19), and which referred to A-replicators—but the confusing and redundant one which refers to L-replicators (Dawkins 1976, p. 30; 1978a, p. 68).

Hull also uses Dawkins’ slogan “longevity, fecundity and fidelity” in another paper (Hull 1980, p. 317), but, on this occasion, citing only (Dawkins 1978a) as the source. Again, Hull goes on to specify that the “relevant longevity concerns the retention of structure through descent” (i.e. L-longevity). Again, he does not comment on the fact that (according to Dawkins) this version of longevity is synonymous with (Dawkins’ version of) fidelity, and is therefore redundant.

The problem is further compounded by Dawkins (1982a, p. 84) where he quotes, at length, from (Hull 1980), and specifically endorses Hull’s interpretation of longevity in this context—thus reinforcing the confusion he himself originated in (Dawkins 1978a).

To summarise, while it seems that Dawkins typically uses “replicator” in the sense of a lineage or L-replicator rather than an actor or A-replicator, he also alternates between the two usages—sometimes with quite bewildering speed. Thus, we have the following two comments (quoted from consecutive paragraphs):

A *germ line replicator* (which may be active or passive) is a replicator that is potentially the ancestor of an indefinitely long line of descendant replicators. . .
[Evidently this refers to actors.]

... But whether it succeeds in practice or not, any germ line replicator is *potentially immortal*. It ‘aspires’ to immortality but in practice is in danger of failing. [Yet now we must be talking about lineages.]

Dawkins (1982a, p. 83)

This confusion between A-replicator and L-replicator is counterpointed (presumably with unconscious irony) by Dawkins’ approving remark that Hull (1980; 1981) is “particularly clear about the logical status of the lineage, and about its distinction from the replicator and the interactor” (Dawkins 1982a, p. 100).

In conclusion: the point of this rather laborious discussion has been to establish that, though Dawkins and Hull make considerable use of the term “replicator”, their usage is quite generally ambiguous as between A-replicator and L-replicator, and calls for very careful interpretation. The recognition of this fact is an essential prerequisite for the analysis of the “unit of selection” controversy which follows.

3 The Elusive *Unit of Selection*

The question I now propose to address is this:

What entities, or kinds of entities, can qualify as the “units of selection” in Darwinian evolution?

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).² This, whatever it is, is the entity which I shall identify as the “unit of selection”.

To make the question even more precise, I am interested in whether the unit of selection is an actor or a 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, a particular kind of lineage.

If “adaptation” evolves at all as a result of Darwinian processes³ then it does so in the form of a change (or accumulation of changes) which, in each case, allows a new *lineage* to selectively displace an old one because it has higher fitness. I take it that there is no substantive dispute about this. Darwinian selection is clearly not a process of “displacement” of one *actor* by another. Similarly, while selection is more usually described in terms of “populations” displacing each other, these so-called “populations” are, in fact, not just any old groupings of actors (no matter how alike), but groupings *which are established by virtue of common ancestry*—which is to say, *lineages*.

²By implication, we are here talking exclusively about “Darwinian adaptations”: that is, “adaptations” brought about by Darwinian evolution.

³There is, of course, never any *guarantee* that Darwinian evolution will lead to increases in “adaptation”—unless one wishes to interpret “adaptation” in such a way as to fall victim to one of the infamous Darwinian tautologies (McMullin 1992b, Sections 3.1, 5.1).

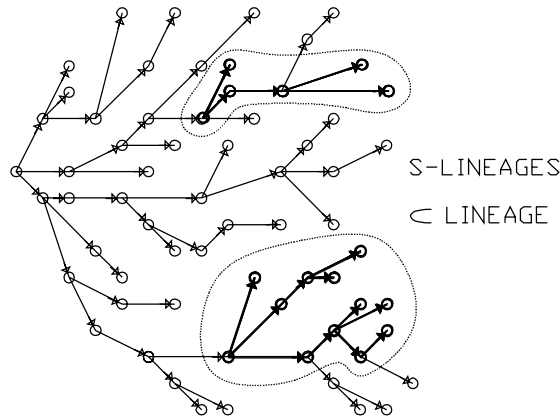


Figure 2: An S-lineage is a subset of a lineage.

By *lineage* here I mean simply the set of *all* descendants of a given *founder* (as already indicated in Fig. 1). Granted, it is only certain *kinds* of lineages which can participate in selection; these are lineages which are characterised by some selectively important attribute (or “trait”). I shall distinguish these by the term *Similarity Lineage* or S-lineage. Briefly, consider some actor which (through mutation or some other mechanism) is different from its parent(s) in some way which is potentially selectively significant. Then I call that actor an *S-founder*—founder of an S-lineage. The S-lineage is the set of all (and only) those descendants of the S-founder which are “similar” to the S-founder in respect of this trait or character (whether “expressed” or just as “carriers”). See Fig. 2.

The S-lineage concept relates to at least some of the common biological definitions of “species”—but can also potentially map onto both higher and lower taxonomic categories! A fundamental distinction from any conventional “species” concept is that S-lineages can, and generally will, *intersect*: a particular individual actor may be a member of numerous, distinct, S-lineages—even S-lineages which are competing with each other. Clearly, the concept is a moderately complex one, which I cannot elaborate fully here. A more detailed discussion is available in (McMullin 1992a, Section 5).

In any case, it follows from this that “adaptations” which are the subject of Darwinian selection must 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 S-lineages are the “units of selection”. However, this needs to be considered in somewhat more depth.

In the simplest case, fitness can be improved by effects which are uniform and consistent for all members of an S-lineage. This kind of effect is particularly easy to envisage, given that all actors in a given S-lineage are already expected to be “similar” to a greater or lesser extent, simply by virtue of their common membership of that putative S-lineage.

Ultimately, this effect would be manifested in one of just two possible ways: if the actors uniformly⁴ live longer than the actors of the competing S-lineage(s), then the mortality (expected death rate) of the S-lineage will be less than its competitors; or if the actors all uniformly have more offspring (per unit time) then the fecundity (expected birth rate) of the S-lineage will be greater than its competitors. In either case, the size of the S-lineage would be expected to increase relative to its competitor(s); and if the difference were large enough, and maintained long enough, the S-lineage would displace its competitors essentially deterministically—which is to say, an episode of Darwinian selective displacement would occur.

This way of thinking leads to associating fitness with individual actors, at least as a form of shorthand. This is always legitimate in a certain abstruse, formal, sense. The members of

⁴Note that I use the word in a probabilistic sense (McMullin 1992a, Section 6.1).

an S-lineage are guaranteed to be “similar” minimally in the sense that they *are* all members of that S-lineage. So any attribute of an S-lineage (such as fitness) can be imputed to the putative “Similarity Class” (S-Class) which also characterises all members of the S-lineage, and thus, implicitly, to single actors, considered as “exemplars” of their S-class.⁵

To put it in a slightly different way, any single actor is necessarily 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 fitness by reference to the activities or attributes of any single “typical” member of the S-class of actors making up a particular S-lineage.

The problem with this point of view is, however, that it is grossly simplistic. There may well be no such thing as a “typical” actor of a given S-lineage. In the simplest case, the selectively important “trait” may not be “expressed” at all in certain members of the S-lineage—but they *can* pass it on, and thus do indeed contribute to the selective dynamics of the S-lineage as a whole. But more generally, it is perfectly possible for the selectively important “trait” to be “expressed” in very different ways by different actors of the S-lineage. Thus, actors of the same S-lineage may be very different in a wide range of characteristics or attributes; or, for that matter, any single actor may engage in a wide variety of different behaviours, or manifest different traits, in different times and circumstances; we require only that these differences not be such as to establish two distinct, competing, lineages with distinct fitnesses relative to the particular selective competition we are studying (otherwise, of course, we would not be dealing with a single coherent S-lineage). The point is that the thing which is selectively important may not be, in any useful sense, a “trait” of a single actor at all, but rather is a “trait” of some groups of actors, or perhaps of the S-lineage as a whole.

Indeed, the fitness of an S-lineage may actually critically rely on the very *distribution of variations* in certain characteristics among the actors making it up. In that case, no single actor, *qua* actor, would even allow a determination of (S-lineage) fitness—except, perhaps, in the extremely contrived sense that we might, in theory, be able to somehow deduce or determine from that one actor what would be the “typical” distribution of variations within an S-lineage which it could found.

We can envisage a range of different ways in which variation in the characteristics of actors 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 actors over these resources or habitats—may well do better than competitors which lack this flexibility (whose actors are more uniform). This can work without any assumption of interaction or mutual recognition between the actors making up a single S-lineage.
2. Alternatively, if there *is* any mechanism whereby an actor can identify other actors with which it shares an S-lineage (with some degree of probability) then there may be other opportunities for *cooperation* between actors, perhaps involving specialisation or division of labour.

In general, both these possibilities could equally be based on dedicating certain 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 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 an actor being so constituted (in structure and/or behaviour) that it functions for the benefit of other actors, at a cost to itself (especially as measured in terms of the number of offspring it produces), is technically referred to as *altruism*, and has historically

⁵I do sincerely regret this further proliferation of prefixes; but it seems to be in the nature of the problems at hand to demand unusually precise vocabulary. If you doubt this, I can only point again at the terminological confusions documented in section 2.

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 an actor can benefit its S-lineage(s) (improve the latter's fitness) *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 actors to other actors in a single lineage) being positively *favoured* by a Darwinian selection process.

I have suggested that altruism could be favoured where there is a mechanism for actors to recognise other actors sharing a single S-lineage. It seems to me that this is a *necessary* condition. By definition, altruism involves costs to the 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 actors which are unconditionally altruistic, and another is otherwise similar (and thus competing) but its 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 *initially* grow at the expense of the original S-lineage of altruistic actors. But recall our basic hypothesis that the altruism of the actors was indeed beneficial to the S-lineage. The new S-lineage, of entirely selfish 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.⁶ 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 actors of a single S-lineage. There are two basic mechanisms which suggest themselves for such recognition:

1. An actor's own "close" relatives (parents, sibling, offspring) are likely to be members (and, to a lesser extent, their close relatives in turn).
2. An actor may attempt to "recognise" other members of some class (which will reliably identify an S-lineage).

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 class and/or recognition of close kin may be quite significant for Darwinian evolution generally. I say this simply on the basis that mutual recognition of this sort is clearly, in the general case, a difficult problem that calls for complex sensorimotor co-ordination, i.e. the utilisation of relatively complex anticipatory models of the world. Benefits (for S-lineages), deriving from more and more effective recognition and discrimination, seem to provide one plausible general basis for a sustained correlation between fitness and knowledge (or "adaptive complexity"), 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

⁶That is, it is not generally the case that such a process would meet all the requirements for selective displacement to go through; in particular, it may not meet the requirement of a *consistent* selective bias (independent of relative S-lineage size).

is the “selective” value *to an actor* of procreation itself? Actually, there is none, and yet everything else that might be an adaptation of an 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 actors, insofar as it has been adapted by Darwinian evolution, as being consistently for the benefit of the S-lineage (procreation of at least some of the constituent actors is an absolutely necessary condition for the continued existence of *any* 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 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 actors, 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 an 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 1978a): the supposed difficulty simply does not arise when one realises that the 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 an 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, in short, is the doctrine of *The Selfish S-lineage*.

4 Interlude: Genic Selectionism

The biological literature on the Units of Selection issue also addresses another question, quite different from the one I have just chosen to answer. This second question might be phrased as follows:

Is there a uniquely distinguished “level” of organisation (biological or otherwise) which characterises entities that can qualify as Darwinian actors?

I suggest that this is precisely the question which Hull set out to answer; whereas it is not clear to me whether Dawkins has ever consciously recognised this second question—not, at least, as a *separate* question. However he does *seem* to suggest an answer (implicitly or otherwise)—namely that *genes* are uniquely distinguished as the only biological entities which should properly be said to be actors. I believe that that claim is quite mistaken (but also that it is questionable whether Dawkins ever really intended to make it).

Specifically, to the extent that we interpret “gene” as a material part of an organism (a fragment of DNA, say), then I accept that such genes *might* be usefully considered as actors (in “suitable” circumstances); but I reject absolutely the idea that they are *uniquely* qualified for this role. I suggest that there are always alternative candidate actors (particularly, but not exclusively, organisms); alternatives which will be *equivalent* in the precise sense of yielding a selection dynamics which is either identical, or differs only by a bijective transformation of the state variables. Preferences among these different candidates therefore arise *only* from pragmatic considerations relating to the particular circumstances in which the theory is being applied.

It is, of course, no accident that “replicators” or “genes” turn up in two quite different senses in the two distinct questions regarding “units” of selection. For Dawkins, genes are the prototypical examples of replicators—and his conflation of both actor and S-lineage into the single term “replicator” is more or less mirrored in his usage of the term “gene”. Again,

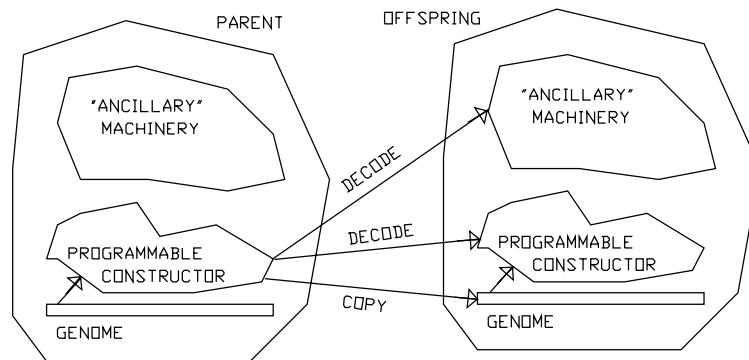


Figure 3: “Genetic” actor organisation.

we should try to distinguish between (at least) actor-genes or A-genes, and lineage-genes or L-genes. The former may (or may not) play the role of actors, and, correspondingly, the latter may (or may not) play the role of S-lineages. Dawkins effectively argues, correctly in my view, that only S-lineages can qualify as units of selection. But he also equates this with a claim that only “genes” (or “replicators”) can qualify as units of selection; I can accept this also, if it is read in the sense of “only L-genes (*as opposed to A-genes*)” or “only L-replicators (*as opposed to A-replicators*)”. But Dawkins seems to go on to parlay this into a claim that only “genes”—now specifically meaning A-genes, as opposed to (say) “organisms”—can qualify as Darwinian *actors*. I believe this to be definitely in error.

To repeat: I reject the idea that a certain level of biological organisation—the “genes” or the “DNA” or the “genetic material” or “replicators”—uniquely fulfills the role of Darwinian actor in terrestrial evolution. I refer to this flawed idea as *genic* selectionism—as opposed to *organismic* selectionism (Wright 1980; McMullin 1992b). My position is the pluralist one that genic and organismic selection, properly viewed, are not competitors or rivals, but merely alternative, formally interchangeable, descriptions of the same underlying biological reality; that, indeed, there may exist an indefinite number of other, similarly equivalent, descriptions; but that, *whichever* of these viewpoints may be adopted, there will be a crucial distinction between actors and S-lineages, with only the latter being properly regarded as units of selection, or entities for whose benefit (Darwinian) adaptations may be said to exist. This latter question, of actor versus S-lineage as the unit of selection, thus cuts at right angles to the question of gene versus organism as Darwinian actor; confounding these two questions leads only to confusion and error.

I may note that I do hold that there is *something* special, in evolutionary terms, about the “gene” level; or, more precisely, about that kind of actor organisation which involves a functional separation into a relatively passive information carrier or “genotype”, and a relatively active information processor or “phenotype” (Fig. 3). This is, of course, the von Neumann architecture for self-reproduction (Von Neumann 1951; 1966a; 1966b), and it exhibits features which are crucially important for the possibility of a sustained Darwinian growth of complexity. Strangely, that this is so, and that this is precisely why von Neumann proposed the architecture he did, is still not generally appreciated even now, almost half a century after his seminal investigations.

Unfortunately, a properly detailed analysis and critique either of Dawkins’ (alleged) genic selectionism, or of von Neumann’s general theory of automaton evolution, would take me much too far afield in the current context. The interested reader (!) can find these discussions fully elaborated in (McMullin 1992c) and (McMullin 1992d, Chapter 4), respectively.

5 So: Does it or Doesn’t it?

So, does the Dawkinsian replicator actually replicate?

Well, in my terms, actors clearly replicate (that is part of their definition). S-lineages, on the other hand, certainly *need not* do so. Of course, to have a *sustained* evolutionary process, new S-lineages have to come regularly into existence; and one can certainly describe this in terms of such new S-lineages being “offspring” of previous S-lineages. In some sense then, perhaps S-lineages *do* “replicate”—but they do *not* “replicate” in Dawkins’ specific technical sense. The whole point of calling something a “new” S-lineage is that it is selectively *different* from, and typically competing with, its precursor(s); whereas, for Dawkins, the very definition of “replication” is that it *preserves* whatever is selectively significant.

Regretfully then, I am forced to the conclusion that Dawkinsian replicators don’t, in fact, replicate! Strange but true.

This is not, of course, the whole story. To make my point here, I too have indulged in simplification. For example, I have glossed over the conceptual difficulties which arise because (under sexual forms of procreation) lineages, including S-lineages, need not be disjoint. Perhaps more seriously, it must be acknowledged that utterly new and distinct *kinds* of actors may sometimes emerge; and that the distinct evolutionary systems which result might interact in significant ways, which would render any explanation phrased in terms of only a single kind of actor quite inadequate; and, indeed, that this *may* already have been a significant phenomenon in biological evolution to date. I have in mind here the kind of hierarchical Darwinian theory described, for example, by Gould (1982), where an S-lineage in one evolutionary framework may actually function as an actor in another. In terms of such an hierarchical theory, my purpose has been restricted to the attempted clarification of Darwinian theory within *one* hierarchical level; but I do not imply, and do not suppose, that such a single level theory exhausts the scope of Darwinism.

6 Conclusion

In summary, my claims are:

1. The “replicator” terminology, as promulgated by Dawkins and Hull, suffers from significant, and confusing, internal contradictions. It would be best abandoned.
2. The *units of selection* in Darwinian evolution are lineages rather than actors.
3. More precisely, the units of selection are *S-lineages*. Unfortunately, these are somewhat complex and even nebulous entities. In particular, they are *not* simply groups of actors which share some overt, inherited, “trait”; nor are they simply “species” in any conventional biological sense. I suspect that S-lineages can only usefully be identified in a *relational* sense, in the context of a particular episode of Darwinian selection. That is, we identify an S-lineage not so much by any similarity among its members, as by the competitively significant difference between two (or more) S-lineages.
4. Actors *may* have a von Neumann-like genetic architecture. This is generally the case for contemporary terrestrial organisms. In this (special) case, an S-lineage can typically be identified or marked by a particular segment of the genetic description, shared by all members of the S-lineage. This makes it tempting to regard this genetic “tag” (or “gene”) as *being* the S-lineage, or to regard this genetic level of description as uniquely preferable for describing Darwinian selection dynamics (“genic selectionism”). *This temptation should be resisted.* A particular genetic tag can only identify a selectively important “trait” *in the context* of the actors which express or translate it; to put it another way, the “meaning” or “decoding” of a particular tag can vary (in principle by an arbitrary amount) from one actor to another, even in the “same” lineage; furthermore, in principle, the same “trait” can be correlated, independently, with multiple distinct genetic “tags”.

5. The logic of Darwinian evolution requires replication with some kind of “heritability” so that coherent S-lineages can form and compete. However, it does *not* require any particular kind of hereditary mechanism, and, in particular, does not require that the component actors have a well-formed von Neumann style genetic architecture.

One reader of an earlier draft of this paper wondered what (if any) pragmatic implications it had for actually designing artificial evolutionary systems. I can only answer for myself of course. The most important implication for me is that, once the mystique of “genic” selectionism is dispelled, we can usefully ask how it is that von-Neumann style genetic organisation can spontaneously come into existence at all. As long as “replicators” (in Dawkins’ narrow sense of fragments of genetic material—or even lineages thereof) are considered as necessary *pre-requisites* for Darwinian evolution, we can hardly even formulate this question, never mind answer it. But once we separate the notion of genetic organisation from the notion of the units-of-selection, we can envisage a form of ALife investigation in which we do not already “wire in” a genetic architecture, but ask how it could arise and be refined—by a process of Darwinian (but not yet “genetic”!) evolution. This is, of course, a fundamentally different direction from “conventional” GA or evolutionary strategy investigations. I suggest that it is a worthwhile alternative to pursue.

Acknowledgements

Perhaps contrary to appearances, it has not been my intention here to be merely polemical. Indeed, I should say that I have the greatest respect for the work of Richard Dawkins—it is precisely because that is so that I have ventured to expend so much attention on it. My hope is that, by expressing my views as clearly, and even trenchantly, as possible, I will facilitate subsequent criticism (and, no doubt, correction) of these views.

This paper arises from 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 indebted to the School of Electronic Engineering in DCU (particularly through the agency of its Head, Charles McCorkell) for continuing encouragement. The final version of the paper has been significantly improved by the incisive comments of the ECAL '95 reviewers, to whom I am very grateful. All errors remain, of course, my own responsibility. Financial support for this work has been generously provided by Expert Associates Limited.

References

- Burks, Arthur W. (ed). 1966. *Theory of Self-Reproducing Automata [by] John von Neumann*. Urbana: University of Illinois Press.
- Dawkins, Richard. 1976. *The Selfish Gene*. Oxford: Oxford University Press. See also Dawkins (1989).
- Dawkins, Richard. 1978a. Replicator Selection and the Extended Phenotype. *Zeitschrift für Tierpsychologie*, **47**, 61–76.
- Dawkins, Richard. 1978b. The Value-judgements of Evolution. *In*: Dempster, M. A. H., & McFarland, D. J. (eds), *Animal Economics*. London: Academic Press. As cited by Dawkins (1978a).
- 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. 1983. Universal Darwinism. *Chap. 20, pages 403–425 of*: Bendall, D. S. (ed), *Evolution from Molecules to Men*. Cambridge: Cambridge University Press.
- 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.
- Gould, Stephen Jay. 1982. Darwinism and the Expansion of Evolutionary Theory. *Science*, **216**(23 April), 380–387.
- 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*. Harvester Studies in Philosophy. Brighton, Sussex: The Harvester Press Limited. General Editor: Margaret A. Boden.
- McMullin, Barry. 1992a (Mar.). *Essays on Darwinism. 1: Ontological Foundations*. Technical Report bmcm9201. School of Electronic Engineering, Dublin City University, Dublin 9, Ireland.
URL (Directory): <ftp://ftp.eeng.dcu.ie/pub/autonomy/bmcm9201/>.
- McMullin, Barry. 1992b (Apr.). *Essays on Darwinism. 2: Organismic Darwinism*. Technical Report bmcm9202. School of Electronic Engineering, Dublin City University, Dublin 9, Ireland.
URL (Directory): <ftp://ftp.eeng.dcu.ie/pub/autonomy/bmcm9202/>.
- McMullin, Barry. 1992c (May). *Essays on Darwinism. 3: Genic and Organismic Selection*. Technical Report bmcm9203. School of Electronic Engineering, Dublin City University, Dublin 9, Ireland.
URL (Directory): <ftp://ftp.eeng.dcu.ie/pub/autonomy/bmcm9203/>.
- McMullin, Finbarr (Barry) Vincent. 1992d. *Artificial Knowledge: An Evolutionary Approach*. Ph.D. thesis, Ollscoil na hÉireann, The National University of Ireland, University College Dublin, Department of Computer Science.
URL (Directory): <ftp://ftp.eeng.dcu.ie/pub/autonomy/bmcm.phd/>.
- Taub, A. H. (ed). 1961. *John von Neumann: Collected Works. Volume V: Design of Computers, Theory of Automata and Numerical Analysis*. Oxford: Pergamon Press.
- von Neumann, John. 1951. The General and Logical Theory of Automata. *Chap. 9, pages 288–328 of*: (Taub 1961). First published 1951 as *pages 1–41 of*: L. Jeffress, A. (ed), *Cerebral Mechanisms in Behavior—The Hixon Symposium*, New York: John Wiley.
- von Neumann, John. 1966a. Theory and Organization of Complicated Automata. *Pages 29–87 (Part One) of*: (Burks 1966). Based on transcripts of lectures delivered at the University of Illinois, in December 1949. Edited for publication by A.W. Burks.
- von Neumann, John. 1966b. The Theory of Automata: Construction, Reproduction, Homogeneity. *Pages 89–250 of*: (Burks 1966). Based on an unfinished manuscript by von Neumann. Edited for publication by A.W. Burks.
- Wright, Sewall. 1980. Genic and Organismic Selection. *Evolution*, **34**(5), 825–843.