

DRAFT

Dynamic Atomism, Fitness Dependent Systems, and The Evolution of Selected Information (Memes) in Humans and Related Species

Greg Laden and Ora Elquist
University of Minnesota

2001

*Please do not cite without permission of the authors.
(Unless in a complementary way, then go ahead.)*

I. Introduction

The present work is a combination of three initially separate efforts:

- 1) Development of a general theory of fitness and replication that does not hinge specifically on either genetic systems or “co-evolution” (which in turn may privilege genetic systems by isolating them as a special milieu for evolutionary process).
- 2) Development of a model for the evolution of certain unique human traits, especially specialization and division of labor (by sex, age, etc.) and increasing dependence on and use of technology for ecological exploitation.
- 3) Investigation of the use of “meme theory” or similar concepts for understanding the relationship between humans and other species (mainly domestic animals).

We have decided to combine these efforts because of similar threads that run through all of them. Each of these investigations considers the role of information that is passed on over time independently of the milieu of information storage or the means of replication. Each investigation recognizes the interdependence of multiple, potentially separate lineages of information. In each case we do (or will in time) consider the possibility that critical information lineages can change their replicative milieu (from one gene to another, from non-genetic to genetic, etc.). For many systems we argue that the change in the form of the system itself, rather than the product of the system, can affect fitness. In other words, what would appear to be the same evolutionary stable strategy (ESS) may be achieved in more than one way, and the different ways in which it is achieved may result in differential costs and benefits.

The following serves both as an “Executive Summary” of our key points and a an exposition of concepts that we feel worthy of further consideration:

- 1) We predict that analysis of archaeological material will become increasingly important to those interested in memes. Heretofore, there have been attempts to bring evolutionary or Darwinian theory

into archaeology (with various degrees of success or interest), and even meme theory has been considered in this context. Here we propose something much simpler but possibly very important: Artifacts as archaeologists define them (the product of human alteration of natural or other culturally produced products) are always, and without exception, the final product of memetic process (assuming there is such a thing as memetic process). This is true because this is how artifacts are defined. Archaeologists place the things they find and study into one of two or three categories, typically artifacts, ecofacts, and “features” (which are either big artifacts or enigmatic stains in the earth). While some products of memetic process may fall into the latter two categories, all things in the first category are such: Therefore, all of the archaeological research done to date has produced a documented record of more memetic process than any large research team could study in several lifetimes. In particular, changes in the nature of the artifactual record over time should reflect changes in memetic process over time. The evolution of the human capacity to manage and pass on information in any way, whether meme theory (or some form of it) stands or falls or is altered in any way, is found in this record.

- 2) We are unhappy with some of the arguments made in support of the importance of meme theory. The idea that grates us most is the argument that memes are real, and important, and even interesting, because they result in behavior that is maladaptive or counter-biological (or bad for your genes, etc.). If memes are real, then they are of interest because they are subject to selection. If so, then they are by definition part of adaptive systems. We do not feel that it would be useful to argue for the relevance of Natural Selection Theory vis-a-vis genetic systems by pointing mainly to really bizarre and hard to explain non-adaptive examples. Yes, genes sometimes seem to have a mind of their own, as may memes, and selection may act at these very reduced levels in ways that startle and amaze us. But for the most part the system of life is interactive and interdependent, so we expect that salient aspects of biological function and evolution will involve interconnectedness and co-evolution more often than not. We do not disregard the importance of “selfishness” of genes, memes, or any other replicator. But we do not feel that the purest (seeming) examples of selfishness are as interesting as the more common interactive systems. Among humans there have been countless instances of reproductive unions and extended families, and there are many adaptive forms of mating and child rearing. By comparison there have been very few monks, and we suspect many have had a hand in reproduction directly or indirectly.
- 3) Memetics has benefited by the transfer of models from genetics. While the power of analogy has been noted, there have been caveats indicating that this may not always be appropriate, because genes and memes are different. We feel that this approach (both the transfer of models from genes and the caveats) is incorrect. A useful and complete theory of replication and fitness should be indifferent at a general level to whether the replicators are memes or genes or something else. Furthermore, there may be things to learn about genes by consideration of memes.
- 4) More importantly, we view adaptive systems as potentially comprised of different kinds of persistent information, some of which being subject to selection. (For a system to be an *adaptive* system, it must include replicated information subject to selection.) Two similar systems may arise that rely on different combinations of replicated, persistent information, with different characteristics of replication. A possible example may be hunting by canids vs. felids. Feral domestic dogs can reinvent wolf-like hunting behavior (with differences) and it seems that canid-style hunting is made up of elements that are essentially “built-in.” By contrast, it appears that felids incorporate important learned elements (cat-memes passed from mother to offspring) in their repertoire of hunting behavior (prey choice, the exact method of killing, and the “idea” that the dispatched prey is food).
- 5) Memetic processes are often attributed greater flexibility than genetic processes. This is parallel to the argument that culture is more locally adaptable than biology, which implies that culture (and/or memetic process) are adaptively flexible. For the latter we note commentary by Sara Hrdy and others, who bring counter examples to bear on this apparent falsehood (such as the “biological” reduction in age of fertility for human females vs. the “culturally” conservative late age of marriage). It may be (and probably is) true that memes are a possible route to local adaptability, but it need not

be the case that all memetically transmitted information serves this purpose. In hominid evolution, it may be that memetic systems are a kludge. Humans are unique because of certain things they do that are distinctly non-ape like (marriage, food sharing, etc.). To reshape a chimpanzee psyche and social orientation to a human-like pattern, powerful forces must be brought to the task, and we would expect memetic adaptations that are involved in this to be highly conservative and relatively inflexible, and probably “stronger” in their influence than more genetic (more innate?) drives and motivations.

- 6) We expect that over evolutionary or long historical time different replicative systems will play key roles in larger integrated systems, and that these individual units may be swapped or replaced over time *within lineages*. As with the canid/felid example, there may be adaptations that are similar on the surface but that involve different underlying replicators. A meme-centered view of the world accords with this idea: Especially if reverse engineering is a means of memetic transmission, one would expect different algorithms or combinations of algorithms to underlie similar final products, while at the same time expecting genetic systems to be more conservative in this regard. This may be true at some time scales, but over longer term this is not as easily demonstrated. Take sex. It is reasonable to expect that all of the metazoans that reproduce sexually arise from a common sexually reproducing ancestor. A gene-conservative view of biology would then predict that there are common elements to the basic and essential process of sexual differentiation (making males vs. females). However, it is known that even at fine taxonomic levels such as worm vs. insect vs. mammal vs. bird vs. reptile there are fundamental differences in how to make males and females, at the genetic level. Does this mean that sexual differentiation among metazoans is a giant and impressive set of homoplasies? Or does it mean that critical elements underlying adaptive systems can swap in and out over evolutionary time? If it is true that different genes and genetically controlled processes can be switched around within a persistent system, then we suggest the following two critical conclusions: i) selection is operating at the level of the larger system at least sometimes, and ii) other forms of persistent information, including those subject to selection (such as memes and genes) can also be swapped in and out of a system. It is intuitive to think that genes are more conservative over the long term than memes, but given that we can only guess about the rates of evolution of either, we should not assume that both sorts of replicators cannot alter in their roles of complex systems over time.
- 7) This should be obvious at this point, but it is worthy of further mention: A model of evolution that effectively addresses these concerns would obviate questions about “nature” vs. “nurture.” If we question the conservatism vs. flexibility of genes and memes respectively, if we allow elements of adaptive systems to change wholesale (thus making old “homoplasies” suspect as such), and if we have a way of describing both memes and genes such that at a general level they are indistinguishable (or differentiated quantitatively rather than qualitatively), then genes and memes become the same thing.

Thinking about Memes: Do we think they exist, if so why (do we think that) and why (do they exist), with notes on how to think about them further.

“Mary” (Subject 17462), from recorded interview 12-22-00:

My morning coffee routine used to involve a French press, meaning I didn't use a coffee maker, but boiled water in a tea pot to infuse coffee in my press. One morning, as usual, I filled the tea pot with water, opened the refrigerator, got out a bagel, put the teapot in the REFRIGERATOR, and turned on the stove. Oblivious to my error, I sat down and stared out the window at traffic, and tried to wake up while waiting for the water to boil, and my bagel to toast. No whistle from the tea pot was forthcoming, so to occupy myself while waiting for the water to boil, I began to read the front page and eat my bagel. Two articles and a 1/2 bagel later, the water still hadn't boiled. This is where it gets really interesting. I got up, OPENED the refrigerator to CHECK on the tea pot (which of course was mysteriously not boiling yet), impatiently closed the refrigerator door, and checked to make sure the stove was on. It was. Mildly bewildered, and shaking my head at the unacceptable amount of time it was taking for my water to boil, I sat back down to wait it out. A short while later, it finally dawned on me: water will NOT boil in the refrigerator. I sheepishly restarted my morning coffee routine thankful that there were no witnesses....

“Hank” (Subject 17463), from recorded interview 12-22-00:

This morning when I tried to make coffee, I poured water all over the top of my closed coffee maker. I looked at the black coffee pot in the dark, shadowed corner of my kitchen, and I saw something that did not look like water in the back of the coffee pot. So I concluded that water had to go into the coffee pot, so I poured water on it. But of course, the back of the coffee pot was closed, and I was looking at a reflection of myself, or the ceiling of the kitchen, or the bottom of the cabinet, or something. It was not water. Therefore I poured water on it. I must have thought "water in back, coffee grounds in basket, pot under hopper, turn on, later there will be coffee"... but something went wrong this time.

“Rulof” (Subject 17463), from recorded interview 12-24-00:

Just this morning, I was thinking about memes while making coffee. I don't know why. I had had a very long night of sleep, and for me that means I wake up physically weakened, and mentally very dull. In fact, I think I might have been thinking about thinking about memes. What could be more mentally dull than that.... Anyway, I was trying to coordinate between having an “early” cup of coffee ... the cup that you grab out of the pot before the water has finished going through.... and dealing with the fact that my filtered water supply was low, so I also was thinking about putting some more water in the back of the pot. Since my “early” cup would be extra strong, I was thinking also about putting extra milk in the cup, so I had poured less coffee in the cup than usual. It was just then that I picked up the quart of 2% milk that I had put out on the counter and poured it into the back of the coffee pot. Disaster.

It is not necessarily the case that we can figure out how something works by observing how it fails under (hopefully) rare conditions, but breakdowns in a complex system can offer clues to underlying mechanisms. In the case of morning coffee brewing, we can guess that the average coffee addict relies on a previously established, habitual routine to eventually obtain a proper dose of caffeine. Morning coffee SNAFU's may be a particularly rich vein to mine for procedural errors (to the extent that test subjects will admit them) if (as seems to be the case) the absence of caffeine in an addict affects the proper functioning of the brain. Indeed, the caffeine conundrum offers an additional twist: Brewing procedures gone awry may indicate not only the form of habitual behaviors by exposing potential errors, but we can hypothesize that the

brain regions that are affected by caffeineation are involved somehow in either the storage, retrieval, or implementation of habitual behaviors in humans.

These anecdotes and others like them suggest several hypotheses. First, that there are stored habitual procedures that can be called up and used under circumstances in which rational on-the-fly problem solving is unlikely to work effectively, or perhaps would be too costly in energy or too time consuming. Second, it would seem that memetic “algorithms” may well exist, because in these cases we see what looks like an algorithm run through with 8 or 9 out of 10 steps in proper place. The fact that parts of an overall procedure can fail and the rest be done as expected suggests the blind following of a script. Third, it appears that such stored habits or algorithms can be called up and used without any significant mentation involved in their evaluation or adjustment: They are simply implemented almost subconsciously, even though they involve activities that under other circumstances (such as the first time an automatic coffee maker is used by a subject) are handled, presumably, with focused, aware, conscious thought.

It is arguable from this that there *are* “memes” and that some of them have to do with making coffee. The algorithm for producing the morning fix of coffee in these two hapless test subjects (who bear little resemblance to the authors of this paper, especially first thing in the morning) is clearly a stored procedure. It is clearly step like, with each step at least somewhat independent of the others. The meme is implemented without immediate reconstruction or engineering, though the test subjects are presumably able to do so under certain circumstances. We presume that the process of making coffee is not innate, but has been learned either from others who know how to do it (by instruction or observation), by trial or error, or in very rare cases by reading the directions on a typical automatic coffee maker (though this phenomenon has not yet been observed by scientists, and probably does not happen very often).

We find this sort of example more satisfying, with respect to the proposal that “memes” exist and should hold our interest as evolutionary scientists, than many other statements about memes. Take for example the often-stated “evidence” that memes exist, that there are memes (such as the meme for being monk) that run counter to biologically sensible strategies. Memes that negatively affect an individual’s reproductive success, but in and of themselves have high fitness, are an important prediction of meme theory, because memes are presumably defined as something that is not passed on through sexual reproduction. However, we are suspicious of any evolutionary theory that does not primarily recognize the fundamental importance of what we will call for now co-evolution of different systems. Imagine a gene that managed to leap from its somatic container and insinuate copies of itself in other individuals. Presume further that the departure from the original host could potentially kill that host, and that in cases where the original host was killed (perhaps by explosion), the gene would spread even farther and more quickly to other potential hosts (by the scattering of little bits of the original host). Virulence (in this case frequent and dramatic explosion of hosts) may evolve but it may not, and more likely, there would be tradeoffs between explosive virulence and subtle sloughing (spreading through skin cells or snot) that would constitute the full range of transmission strategies for this gene. We do not focus on, nor do we exclusively care about, the most virulent strategies (the exploding hosts, or even the devout but somehow charismatic monks). We are concerned with the whole process, the strategy of replicators with respect to the longevity, or lack thereof, of a host, and the rate of transmission. (Obviously, our exploding victim struck with the jumping gene is meant to be analogous to pathogenic infection, see: Levin 1996, 1997)

Of possible approaches to study of and thinking about memes, the most important comes with the idea itself and has been espoused and implemented by Dawkins. Simply put, if memes exist, then they are Darwinian (virtually by definition), and thus analogies with another probably better understood process – genetic replication, transmission, and expression – can be useful. This we think is good, but we also note the following. First, it is very rare these days in the discourse on genetics to switch more or less randomly between discussion of sequences of coding DNA (genes/alleles) and their products (proteins) or their final effects (other aspects of the phenotype) as though they were all exactly the same thing. This, however, is virtually ubiquitous in the discussion of memes. A baseball bat, the way to swing a baseball bat, the desire to watch somebody who is really good at it swinging their baseball bat, and the actual swinging of a baseball bat are not all the same thing, but they might all be called memes, and worse, they might all be called the same meme. We are also concerned with the wild meanderings through levels of selection, or more exactly,

where in the heterarchic complexity of life system one should (or can) measure fitness, that seems to happen in this discussion. Related to this, and parallel to the counter-biological issue, we are concerned about the general principle that memes *must* exist and *must* be Darwinian because there are ideas that most rationalists would think of as counterproductive and stupid that seem to spread like wildfire. An adaptations argument about genetic systems that focused on the spread of really dumb “adaptations” would not be very convincing.

The flavor of meme-related discourse seems to have this aspect: Memes are both everywhere and everything, or there are no memes at all. Either information and knowledge is particularized, analyzable as fitness-differentiated entities, or everything is learned, stored, and processed by a general problem solving mechanism – or a set of modules – in an evolved brain. One of us remembers those “heady days” back at an east coast university hanging around with one of the world’s leading evolutionary theorists who was also a very close friend. We would habitually hang out in a certain lounge in the Anthropology department, eat lunch and/or dinner together every day at a local grill, taught together, etc. We came, together, to see the world as one unified and continuous fitness landscape. One of us would observe some oddity in this fabric of differential fitness, and glance over at the other, and we would both know that we shared a common amusement at the same wrinkle in $W = I + D$ (or perhaps $c < br$). To us, memes were a natural extension of the presumption of the time-averaged effects of the environment on the fitness of genes. Those were the days of the first computer viruses and the first email chain letters. We knew that replicating systems and fitness models were leaking out of wherever they existed into the electronic world, and this seemed so natural and inevitable that we never even questioned it, nor did we seem to even care about it. Memes were not our interest at the time, so we let the idea of Memes slide into our thinking without much thinking about it. After all, the idea was proposed by Dawkins, and he seems a smart guy.

Eventually, however, one must “get real” about shifting our day-to-day amusements into the framework of scientific scrutiny. Even the most reductionist biologist cannot look around at a world of phenotypes and see only genes. Nor can the most deterministic geneticist see every biological thing as unconnected with non-genetic factors. And, even the most well adjusted adaptationist cannot see absolutely everything as an adaptation (though perhaps nearly so). It is quite possible that there is something going on in the middle: There are memes, but not everything is memetic.

We have some of our own recommendations about how to think about memes. First, we suggest the inclusion of all of the standard biological concepts, such as analysis of fitness, costs, benefits, etc.; resource limitations; population level systems (see Boyd, etc.); environmental context and norm of reaction; and Romer’s rule. Pursuant to this, we think considerable time must be spent examining Darwinian replicator and fitness theory, and not focusing on memes in particular. Indeed, by separating a replicator from the particularistic and historically strange system of replication and transmission by DNA molecules, we believe that we can gain insight on Darwinian process that may feed back on more traditional questions, especially having to do with localizing and measuring fitness, and levels of selection.

One of the consequences of examining replicating systems independently of their physical context (as genes, memes, whatever) is this: We propose that fitness operates at local levels that are only very rarely individual alleles (or genes) or memes, but rather, at a more dynamically defined level that almost always involves more than one “element” (which may in turn be an allele or a meme). Fitness, therefore, can and perhaps usually does exist as a property of combinations of things, including replicators and non-replicators. We find it hard to imagine an organic system that includes memetic process whereby memes would not be part of the overall function. Memes, genes, abiotic factors, and other elements interact with a certain degree of indifference as to the replicative milieu of their (sometimes) co-evolving partners in fitness enhancement strategies. For this reason, we feel uncomfortable with the memetic drive idea as it seems to be asserted by Blackmore. Yes, there can be “drive” from one part of a system or another, and certainly genes and perhaps memes (if they exist) can appear very strategic, and will not relinquish a selfish position in any model, even an interactive co-evolutionary model with multiple levels of fitness. But we feel that the full-blown memetic drive model is overdone, and carries too much burden both in demonstrating that memes “must” exist (because monks exist) and as a model of how this putative replicative system operates. We know of no genes that operate in isolation ... every single gene requires the proper functioning of other genes. We know

of many genes that require non-genetic factors to operate. And, we suspect that memes, if real, are not much different.¹

¹ A note on “Evolutionary Pattern” and memes: The question of emergence of memes has at least three levels. First, if memes are a feature of a particular kind of organism (such as a human with a symbolic mind), then when, under what circumstances, where, and how did this capacity arise? Second, how does a particular meme arise? To this we would add the idea that since memes can (to the extent that they exist at all) exist outside of individuals, and that there may be more sophisticated and complex things that are sometimes called “memeplexes,” that we can sometimes treat memes (or memeplexes) as a sort of species. So, the third level of question may be: Do memes or memeplexes speciate? Is there a pattern of meme diversification, radiation, or even extinction?

Within this context lies one of the possible examples of inappropriately driving details of meme theory with presumptions about genetics. For now we’ll use the term “memeplex”. It could be said that a memeplex may arise from the combination of different memes, themselves possibly part of previously existing memeplexes. This may be said to be a difference between memes and genes, because we don’t normally think of organic speciation as involving the recombination of genotypes between species to make up a new species. Thus recombination of memeplexes to make new memeplexes may be a special feature of memes. However, it is not the case (of course) that organic (genetic) evolution is different from this presumed memetic pattern of evolution. Hybridization and symbiosis of bacteria and primordial animal or plant cells (to give rise to eukaryotes) are examples of “recombination” of organisms to generate new organisms.

The archaeological record does seem to give examples of meme radiation, diversification, and even extinction. Pottery making arises around the world at roughly the same time (in the big view ... the origin of most pottery traditions can be traced to a period of time that is only a few thousands of years long across the globe), and in at least one instance – Polynesia – pottery making ceased. To the extent that pottery production and/or use is facilitated by memetic process, we can see in this pattern many of the aspects we attribute to organic (genetic) macroevolution: Possible homoplasy (that arises because conditions are similar around the world?), speciation, radiation, diversification (of both form and function of pottery) and extinction (though not mass extinction in this case).

The reason why this is important is simply this we wish to penetrate the problem of memetics with a more general theory (or at least description) of replicating systems (following Dawkins and Dennet), and any such theory should (we hope) inform us of both macro evolutionary pattern and micro evolutionary processes. In addition, macro evolutionary patterns are potentially very interesting with respect to memetic evolution, especially with reference to the evolution of sociality (and possibly eusociality), including for humans the evolution of domestication, complex societies, and eventually a world system.

The Basic Fitness Dependent System: The Smallest Locality of Fitness and Dynamic Atomism

Here we begin the development of a formal model for the function and evolution of replicators, such as genes and memes, that allows for the systematic characterization and consideration of the following:

- The location, or unit of selection (the physical or action related locus of variation in fitness)
- Contingency (local context-dependent variability in the product of expression)
- Incorporation of genes and memes and similar constructs into a single model
- Development of a system-level description of processes allowing one to avoid the 'nature vs. nurture' dichotomy (?)
- The proposal that "system level transition" is an important evolutionary process

We propose that the primary (direct) location of selection is not the replicator (i.e. gene/allele), but rather a unit called here "element" which is a contextualized product of a replicator (i.e., a protein) and essentially related products. There are, however, important differences between the concept of 'element' and 'product.' Contingency refers to the local variation in what happens when a replicator is expressed, as an effect of context. The appearance of a product might depend on interaction with other elements at the time of expression, or its exact function may depend on which other product it interacts with. Contingency is a critical aspect of replicators, and in fact, its importance is the reason why we suggest moving the minimal unit of selection from the replicator to a slightly higher level (element).

We propose that the function of replicators is understood only in the context of a system. A system may be a metabolic system in which the contingency shaped product of replicators such as genes interacts over time. This calls for a third aspect of replicators, other than replication or expression (the usual two aspects referred to), which we call "*system role*". The system role of a replicator is defined most simply as the position of the product of the replicator as an *element* in an interactive system of any kind. The system role is the first order phenotypic realization of the product, and the actual minimal unit of selection for the replicator.

A group of elements that interact forms a system. In some cases a system may have the property of coherency or completeness in that it functions in some identifiable role with all the elements present, and can be modeled or conceptualized separately as it is represented as a whole. A system that is complete in this sense, and which would lose this property of completeness if even a single element is removed, is an "atomistic" system. It is probably not important at this time that this definition is a bit slippery and hard to use, however it will be handy to consider the possibility of notationally representing sets of elements in a system as a single symbol, and in doing so attempt to only represent systems that are reasonably coherent and complete in this way.

We define two kinds of elements. Replicator based elements (R-elements) and Non-replicator based elements (N-elements). R-elements are generated by expression of replicators, (genes and memes for example), and N-elements exist by other means. Examples of N-elements include gravity, water, naturally occurring minerals, etc.

A large part of the model presented here will have to do with how elements are defined and constructed. For now we simply propose that the minimal immediate (direct) site, or unit of selection is not the replicator (algorithm) or the phenotype, but the element itself. We also note at this time that in defining a system of elements, an atomistic system can serve as an element, and any given atomistic system can be considered as a unit of selection.

Perhaps the easiest, and most convincing, way to think about this is to consider the relative fitness of two alleles (a and A) of a single gene in a very simple organism. In this organism, all alleles of all genes utilize a variety of amino acids, but none uses a particular amino acid (let's call it "X" to avoid arguments about how important someone's favorite amino acid may be). But there is one exception, in our focal gene,

Allele A uses amino acid X, and this is the only instance of the use of this component. Further assume that the product of Alleles A and a are functionally identical in every possible respect. From this we would assume that $W(a) = W(b)$.

But wait, there is more. In this particular example, the tRNA for Amino X is coded for by a faulty gene, and it does not work at all. Therefore, even though the product of the algorithm represented by alleles A and a should be identical, and by all reductionist appearances A and a have the same fitness, $W(a) > W(A)$. But because of the non-functioning X tRNA, allele A produces no useful product. From this example we can see that the proper functioning of a given replicator depends on the auxiliary parts of the system of expression (and transmission as well). Similarly, fitness of a given replicator is affected by other co-products. The atomistic unit in this case must be minimally Allele A and all of the genes that code for the relevant tRNA's, not to mention ribosome's, mRNA's, etc. etc. If anything goes terribly wrong in any one of these elements, the allele is ineffective and has neutral fitness or is selected against (depending, further, on its system role). From this it can also be generalized that the fitness of a replicator goes down on average when it is involved in a larger number of relationships.

In this model selection does not act on the gene or meme itself, but on the element, which is defined as the replicator, its auxiliary parts for replication and transmission, and critical co-products, and indeed, more subtly, selection actually works minimally on a particular element further defined by its system role.

The incorporation of genes and memes as well as other replicators (should they exist) and N-elements into a single scheme has the advantage of allowing us to model complex systems that rely in part, perhaps in large part, on replicator elements that are subject to selection. Selection as a contingent outcome of special conditions (see Endler) is certainly blind to the *modus operandus* of the replicator itself.

System level transitions: Characteristics, costs, and benefits

System level transition is the change (among the organized units of one or more systems) of the *loci of elements in relation to each other*. A simple example is the movement in many taxa (most notably the Order Primates), of the element in vertebrate physiology that produces ascorbic acid from within cells (where a gene codes for an enzyme that is critical for synthesis of this vitamin) to outside the organism, where the product is secondarily obtained through ingestion (i.e. of fruit). This particular evolutionary event, basal to and thus derived for the Order Primates, may have been a random genetic event combined with the good luck to have sufficient quantities of Vitamin C in the diet. But, it has consequences having to do with the reliability of the system by forcing a degree of 'specialization' (what if there are insufficient supplies of Vitamin C?) and efficiency (while there are natural supplies, the organism does not have to divert energy into production of the nutrient).

The most interesting and probably the most important system level transitions involve systems that are atomistic with respect to a particular function, and that can thus be defined themselves as elements. The vitamin C example is a case of this: Vitamin C is a critical nutrient, so the system for delivery of this nutrient to its location of activity, the location of activity itself, and the production of Vitamin C are all part of a single element in the physiology of an organism. If modeled at only a modest level of generality, the Vitamin C-centered system is identical for an antelope that synthesizes it and for a primate that must ingest a plant that synthesizes it. However, these two systems are critically different in that one involves cells in a single organism, the other multiple tissues spread across multiple organisms. Analysis of this transition should involve measurement of the cost and benefit differences between the two forms. Additionally, the two systems are different in the opportunities or limitations that arise for further evolutionary change: A primate will not do well in an evolutionary sense in an environment poor in vitamin C. Some humans, finding themselves in this situation, suffered a cost, and subsequently developed (memetically operated) means of dealing with that cost (an economic solution involving the exchange, storage, and transport of fruits).

In the same example, we also see potential for subsequent evolutionary consequences in the possible added attraction some fruits have for primates (because of their vitamin C content), thus allowing further exploitation of primates as seed dispersers by plants

Here we are also able to invoke Romer's rule: The primate is just trying to get its vitamin C, so it eats fruit. The relatively complex economic/dietary practices developed to keep sailors from getting scurvy are an extension of the same. The most generally modded form of the Vitamin C-centered system remains the same, but more and more stuff has to happen to make it work, and most interestingly, involving more and more individuals, increasingly specialized to task, across multiple species, with a great deal of cost expended.

Below we will propose a similar kind of system transition related to the origins of key modern human foraging behaviors as well as for later developments in the production of food (plant and animal husbandry). But first a few more words about the generalized conceptualization of replicators.

Replicators

Replicators are entities that are subject to Natural Selection through the interface of their products with the environment. What is usually called a meme is a form of a more general class of replicable algorithm. The algorithm is information that underlies a process, and is similar to a gene in that both are algorithms for the sequence of steps to produce a product.

Following the genetic analogy, we say that a meme is "expressed" when the algorithm is carried out, and that there is a product, which may be a physical thing or a series of actions or activities. Furthermore, there may be differences in the final product depending on context, thus integrating the idea of "contingency" in meme functioning. It is probably true that not all "ideas", "concepts", or "symbols" are memes. The broad brushed application of the idea of memes to anything that has to do with a human mind is probably counterproductive, and minimally, annoying. This approach also limits the use of the replicator concept of a meme to things stored in the human mind, which for our purposes here would be very counterproductive.

Furthermore, the thing that we call a "meme" is not the product itself, any more than a protein is a gene. A ceramic pot is not a meme, though the process of making a pot, which is an algorithm or set of algorithms, is. The pot is the product of the meme's expression.

It is possible that, while genes and memes are "replicators", that there are other kinds of replicators that do not fall easily into either of these categories, which is a subject we will not take up here. Generally speaking, however, we propose here the possibility that there can be more than one kind of replicating algorithm. As long as the entity in question meets certain conditions, it can be considered as a replicator. Though the mode of replication and inheritance for memes vs. genes is arguably different, the outcomes of these different modes still allow the process of Natural Selection to apply.

The Three Necessary and Sufficient Conditions of Natural Selection are sort of an acid test for traits shaped by selective forces. [1) variation in a trait, 2) inheritance of a trait, and 3) differential reproductive (replicative) success.].

Variants of genetic or memetic algorithms for a product must be differentiated from variants in phenotypic products produced by identical algorithms. Variants of algorithms will be expected to have differential fitness in selective environments, and can thus be observed to differ in frequency in a population of individuals in time and space. Variants in phenotypic products produced by identical algorithm are due to contingency in **expression** (not necessarily related to system role of an element), but the range of variation will be kept in check in a situation where selection is occurring and through fidelity of replication.

Novel variation in genes occurs through mutation, where a random error occurs, for example, in the replication of DNA, and is only important in terms of selection when the mutation is inherited by offspring. Much of the "new" variation inherited and expressed in an organism is the result of novel combinations of alleles. Mutation may often produce no visible effects because of the redundancy of DNA's coding system whereby the same amino acid or product can be produced by variant codons. In a way, this is sort of an "independent invention" of the same phenotypic product. One can extend this argument to memetics.

However, the distinction between new memetic variants produced by mutation vs. invention first needs to be addressed.

As with genes, mutation in the replication of a memetic algorithm can occur as a random error during replication, which from the meme's point of view as opposed to the organism's point of view can be deleterious, neutral or beneficial (as in the case of serendipitous moments). Only in circumstances of transmission will selection be important. As memetic algorithms are information on how to perform a series of steps for the production of a product, mutations can occur through addition of a step, deletion (pouring water onto the back of the coffee maker without first lifting the lid), translocation (trying to boil water in the frig vs. the stove), etc. (Clearly, most novelties are deleterious.) The invention of a novel idea can be considered a kind of mutation, but perhaps most ideas are best thought of as "recombination". Ideas typically do not occur in a vacuum, and "lightbulb" moments are mostly the products of mental abduction where information is brought together in a new way.

Inheritance of a trait, or transmission must occur in replicating algorithms, though the *mode* or means of transmission is probably irrelevant to the process of selection if selection occurs at the level of the phenotypic product or element (i.e. trait). The important thing is that if transmission occurs with a high enough degree of fidelity, then longevity of the algorithm is preserved. However, there are key differences between genetic and memetic algorithms worth discussing. These differences mainly lie in each kind of algorithm's mode of inheritance or transmission. Genetic replicators entail a 'digitalized' means of replication and transmission, but this characteristic is perhaps less certain for memetic algorithms. However, both replicator systems possibly utilize analogous if not identical 'reproductive' strategies. Memetic transmission clearly differs from genetic transmission at least in terms of genealogical lineage. The differences may be briefly summarized as follows:

Genes transmit from a parent generation to an offspring generation, and thus transmission is vertical. Memetic transmission involves bequeathing the algorithm from one generation to the next, however this process is not limited to vertical transmission, but can be horizontal/lateral as well resulting in different 'classes' of generations. The main difference this would produce between the two different types of replicators is the rate at which evolution can occur. This means that the process of memetic transmission could be described as Lamarckian in that memetic algorithms are 'acquired'. This is because memes are transmitted through learning, observation, imitation, and probably to some extent, reverse engineering. While this means that memes may be transmitted in a Lamarckian fashion, it does not make them necessarily subject to Lamarckian evolution to the extent that Lamarckian evolution incorporates the idea of intentional striving of an organism to improve some characteristic of itself.

First, memes do not evolve for the benefit of the organism in whose brain (or wherever) they reside (though advantages to the organism can and often do occur as a side benefit), but memes increase in frequency in and of themselves due to some characteristic advantageous to the meme itself. Second, memes do not consciously strive for self betterment (although perhaps there could be a "self improvement" meme, or a meme for what is perceived a better or worse - bigger is better, etc.), or for greater efficiency but instead are subject to the "blind" effects of selection. Therefore a memetic variant will become more common at the expense of other competing memes based on random environmental circumstances. In light of these points, memetic algorithms can only be subject to Darwinian evolution and not Lamarckian evolution.

Differential replication is the means by which the end result of selection can be observed. Those replicator algorithms that produce a phenotypic product with higher fitness will become more common through time at the expense of it's less fit competitors. This can be observed by looking at the most frequent/optimal variant representing the mode in a population, and this is perhaps the best way to examine the evolution of memetic algorithms. We assume here that memetic replicators do exist, and take as evidence the reality of non-genetically manifested products where the mode functionally represents the optimal trait variant. Artifacts in the form of pottery can be used as an example of this (Neff, 1992). Although Neff does not explicitly state his perspective on ceramics evolution from the perspective of memetics, he does say this:

My basic argument is that ceramic evolution is differential persistence of *information* on ceramic making through history. A theory of ceramic evolution is therefore a theory of how information on pottery making is invented, transmitted, recombined, and eventually lost. I suggest that the mix of ceramic traits observable at any point in time-space results from differential persistence of information within ceramic traditions and from differential persistence of traditions *themselves* [emphasis added: Neff, 1992].

An example of this differential persistence or replicative success is presented by Neff (after Braun, 1983). For Woodland period pottery there was one overall long-term trend of an increasing frequency of pottery vessels with thinner walls through time. Functionally, this was attributed to a subsistence change towards seed that required longer periods of boiling at higher temperatures. In functional terms, thinner walls were more resistant to heat stress and increased thermal conductivity. In this case, thinner walled vessels were the most optimal trait variant in the overall ceramic population.

Generally, there are considered to be three Necessary and Sufficient Conditions of Natural Selection as outlined above. We (and others such as Lewontin?) argue that there is a "4th" Condition: that of fidelity. Fidelity is a necessary attribute of the replicating algorithm in that algorithms are codes for the manifestation of some product, and if the algorithm varies, the resulting product will also vary resulting in unpredictable variation. If the algorithm did not replicate with a high degree of fidelity, then the variation in the product or trait produced by the algorithm would be 'selected' for/against in a random manner, and the resulting 'algorithm pool' would exhibit a pattern more akin to random drift than one affected by forces of Natural Selection, and Natural Selection is a *non-random* process. Therefore, Natural Selection cannot occur without fidelity in algorithm replication.

Fidelity is a property of a replicator whereby some threshold of frequency of exact copying during replication is achieved. Anything below this theoretical threshold excludes the entity from being a replicator. Fidelity is inextricably tied to longevity of an algorithm as well, for without fidelity, longevity is impossible. The digital like system of genetic replication clearly confers the property of fidelity onto genetic algorithms, but can this same property be applied to memes? [Note: The authors are not decided, nor necessarily do they agree, if there is a "threshold" for fidelity.]

Language is certainly a system of digital like code by which a meme can replicate, but this alone does not lend itself to understanding how memes are replicated through observation, reverse engineering, inference or imitation. One could expect the costs of each of these means of meme replication to vary. We can assume here that the less 'digital' a means of replication is, the more costly it is to replicate the algorithm accurately. Since there must be a certain level of fidelity for a replicator to even exist, some mechanism must exist to increase this quality to allow for meme replication.

One such means for a meme to achieve the critical threshold of fidelity would be repetition. It is interesting to note that oral histories and stories from cultures without a written language are often characterized by frequent repetition. Practice makes perfect. Repetition is important for habit formation and memorization of information/code for steps to produce some product. Most importantly, over time repetition would serve to reduce the variation in the product produced by the memetic algorithm. In this sense, memetic algorithms may not have a true digital like system of replication the way genetic ones do, but repetition would allow memes to have some semblance of this - one effective enough to reach a reasonable level of fidelity. In fact, repetition may allow memes a certain advantage over genes in this respect. Most of the time an error in genetic algorithms is deleterious. Genetic algorithms can't afford to make fatal errors, and repair mechanisms do exist in genetic systems to reduce this possibility. In terms of memes, however, not correctly playing Fur Elise on the piano the first time won't kill you (excluding some twisted circumstance). More importantly, it won't necessarily kill off the Fur Elise meme either. Products produced by memetic algorithms get to have a learning curve or mechanism of self-correction perhaps analogous to, and maybe more forgiving than, the repair mechanisms extant in genetic systems.

Memes can be transmitted through observation, inference, reverse engineering, or imitation of making a product without the benefit of language (although it certainly helps). A simple example might be watching someone tie a square knot for the first time. What's being learned by the nascent knot tyer is not

how to hold the rope in their hand, or the pulling of the rope, but where to hold the rope and when to pull on it. The assumption made here is that there is a basic set of algorithms for moving body parts that we all acquire/have early on. The new meme consisting of where and when to manipulate the rope is probably engaging these basic sets of instructions of how to move certain body parts. Perhaps an essential part of memetic algorithms is an *expectation* of the correct result. When one doesn't get this, one simply repeats making adjustments until the expected result is achieved. What this means is that the body being used by the meme to produce some product has to practice how to carry out the meme's instructions, suggesting perhaps that memes do have the property of fidelity, even if the organism does not in carrying out its instructions. Fortunately for the memes, practice reduces the variation in the product expressed making it subject to selective forces.

This means that the initial cost of acquiring or rather, transmitting the meme is high as many potential errors in the steps coded for by the algorithm could occur, but over time, the cost is significantly reduced. Fidelity in the product is also increased, as one would get less variable results in the product expressed. Variability will inevitably occur, either through mutation or variable expression, but one could expect that this variability would be far less common in highly functional products as opposed to 'stylistic' ones. This suggests that the threshold of fidelity is higher in functional memes. In terms of selection and differential replicative success, an environmental factor might be "intentional" cost-benefit analysis, or judgments that people make which is in turn based on their own cultural memetic milieu, and direct/indirect experience which fuels foresight.

Furthermore, memes probably use strategies to increase fidelity as those memes with higher fidelity will have an advantage over those with less fidelity. While the "digital" nature of memetic replication cannot be demonstrated at the moment in say, the human brain, one can address writing systems as means of faithful replication. The meme(s) for a writing system might be appropriated by any meme as a strategy on the part of the memetic algorithm to first increase the likelihood of fidelity, and secondly, to increase the likelihood of transmission. Both are interrelated. In regards to the first proposition, once something has been "carved in stone" so to speak, replicating it through direct observation or rather reverse engineering perhaps becomes less costly than through memorization (disregarding habit for the time being). Secondly, the physical production of writing onto a material provides *external storage* of a memetic algorithm. Writing now allows information to reside somewhere else other than a brain. This increases the likelihood of transmission, because it reduces the cost of the reverse engineering involved in "reading someone else's mind", filtering out someone else's errors in transmission or phenotypic expression of the instructions, and increases the meme's capacity for diffusion by increasing its 'reproductive opportunities'.

Transitions

Symbolizing, language, ecology, and the relationship of the human mind to the meme:

How does a “symbolic” mind fit in with a theory about memes? It has been proposed (Blackmore, for instance) that language evolved for the purpose of spreading memes, and in this regard, Blackmore even goes so far as to imply (or even state) that the evolution of language is roughly independent from other human adaptations: Human language is a byproduct of memetic drive. We take issue, as we have above, with the proposal that we should think, let alone be enamored by the assertion, that memetic systems have evolved as purely selfish independent systems despite the fact that memetic fitness cannot be separated from other localities of fitness (such as genetic fitness). It is possible, following Deacon, that language is a product of memetic systems, and it is possible that memetic systems are a product of the existence of language. However, we would be surprised if either was true, and we suggest that memetic systems and language capability are a) intertwined in their evolution, ontogeny, and functioning and b) neither system is exclusively involved with the other. Both systems (to the extent that they can be separated) are likely rather promiscuous in their associations.

For the time being, we will assume nothing about the association between language capacity and symbolic capacity. This is partly because we do not pretend to be able to define, measure, or scientifically observe either, and partly because such an assertion (one way or another) may be irrelevant to the arguments we make here. Having said this, we wish to propose here, briefly, a model for thinking about symbolizing (the current fashionable term for this process, whatever it turns out to be), and then to suggest how memetic process and symbolizing may be interrelated, at least in humans. We quickly add, however, that we do not wish to restrict memetic process to symbolic minds. This weakness of will (or great insight) on our part will require us to refer to “symbolically mediated memetics” (SMM) because we do, despite our reservations, believe that memes generated, stored, and operationalized in a symbolic mind are different from other memes (such as those in NH animals), and this difference leads to some very interesting hypotheses about the interaction between humans and animals (and even plants) such as the process of domestication.

Terry Deacon once made this remark about symbolic capacity: If all you have is a hammer, you will treat everything like nail (presumably this metaphor originated in a fortune cookie but was put by Deacon to good use). His point was that we cannot distinguish between human mentation that is symbolically dependant and that which is not, because we have the hammer – symbolic capacity – that we use for everything even if other organism or some of our ancestors did not. This is an important caution when thinking about memes in a symbolic species: We may see memetic process in humans as always symbolic because it probably is always symbolic in this context, but need not be, and need not have been in the past.

There are a lot of ideas about how and why a symbolic mind arose, including Deacon’s idea that it had to do with making social contracts (marriage) that would be difficult for an ape-like ancestor. Deacon associated this idea with hunting. Wrangham et al (including one of the present authors) have modified this idea, suggesting that the controlled use of fire, integrated with central place foraging (necessary if fire/cooking is used in food processing) selected for agreements among early *Homo erectus* to not act like a bunch of apes and get along better. Laden has further suggested that investment in technology such as handaxes, which cannot be protected in a chimp-like dominance system, either selected for or at least indicates the existence of (via the archaeological record) similar kinds of social contracts, even if in a nascent form. All of these interrelated ideas are associated with one critical point: Symbolism (or symbolic capacity) allows the mediation of social relationships of a type that are unthinkable in a chimp-like mind. Symbolic capacity is an evolutionary kludge to turn a chimpanzee-like ancestor into a latter day hominid (and the debate as to how much “latter” will continue for some time).

From this point of view, it could be said that memetic capacities in humans are enhanced and potentiated via symbolic process and associated linguistic abilities. This is bad news for meme-jingoists and

aficionados, because it appears to make memes a side effect of symbolic/linguistic evolution, rather than the other way around. However, we are proposing something different and we think more interesting.

The symbolically mediated social contract between individuals that we see in pair-bonding/marriage, and in the kinds of foraging activities we see in modern hunter-gatherers (central place foraging, food sharing, specialization and division of labor, investment in technology, and complex territories) involves an evolutionary novelty (for apes): *more than one individual shares a role in an adaptive system that can be defined as atomistic with respect to fitness.*

This is not about group selection. The groups of foragers who practice one technique or another for extracting energy from the environment and turning into babies are not the unit of selection. The unit of selection is the ecologically oriented strategy of specialization – which probably often increases risk of success – and subsequent sharing (to reduce variance and make for a whole dietary system). Specialization and interaction between specialists is a kind of domestication among the domesticators. A human herder may use this and that meme from a herd animal in combination with this or that meme from a domestic dog with his/her own memes (for gaining access to animal products) to produce the same thing that a forager forbearer did (thus a herder is a forager desperately trying to remain a herder). Similarly, a group of foragers specializing in their efforts and sharing some or all of the products over time acts like a meta-individual who happens to be a very good generalist.

Domestication and the evolution of agriculture

We can apply the idea of replicator systems to the example of domestication and the evolution of agriculture. In a sense we can view the switch from hunter-gatherers to farmer as a shift from a foraging "system" to an agricultural "system". Agriculture is the outcome of a co-evolutionary system between people and their domesticated plants and/or animals resulting in mutualistic relationships. The issue of intention in domestication as defined by Rindos, (1984) is irrelevant and perhaps non-existent, though on the surface of things intention must play some role in situations where agriculture is adopted particularly as it concerns expectancy.

Humans did not 'invent' agriculture. That distinction probably belongs to ants who have been practicing agriculture for about 50 mya. However, there is an implicit distinction made between ants and humans in terms of agriculture. The agricultural system has evolved as a symbiotic relationship between ants and their 'domesticates' where each species involved is in a mutually beneficial exploitative relationship. At worst, human agriculture is often presented as people somehow having imposed their will upon nature, and at best, that agriculture is a human adaptation where humans *select* for desirable traits in the domesticate, or as an adaptation to various environmental changes such as resource stress, or environmental shifts, yet the process of domestication as an adaptation is still viewed as a one way street, and so are incomplete..

Other explanations for human agriculture that have emphasized social factors, such as competitive feasting (Hayden, 19??) reify implicit assumptions about humans and nature. Social factors typically are not considered part of the "environment" in which Natural Selection operates. This is where memetic algorithms have the most to contribute. The fact of the matter is, humans are not independent of their environment.

Rindos takes an evolutionary view from a cultural selectionist perspective in that agriculture (and its inherent information which we would class here as a set of agricultural memes) confers reproductive success on the organism practicing it. Agriculture as a subsistence strategy certainly can instigate population growth, but one can just as easily argue that it conferred replicative success to the agricultural meme(s) itself. How can this be justified? First memes are not limited to biologically 'non-adaptive' behaviors. It would also be ridiculous to suggest that agricultural behavior is genetically programmed. Secondly, and perhaps most importantly, the key here is expectancy. As stated previously, expectancy of a certain outcome is an integral part of any meme. Simply put, there is a right outcome and all other outcomes are wrong.

If agriculture consists of a set of memes, they are highly 'functional' ones in light of the fact that the product pertains directly to the food supply of the organism replicating them (this applies to all agricultural material and non-material technology). One possible expectancy that would certainly make sense would be a plentiful food supply, or in more abstract terms, an increase in carrying capacity (though it's possible that

there might be other expectant outcomes). Those memes that best meet this expectation will have greater replicative success. Therefore, it should not be surprising that that high memetic fitness in this scenario would also 'piggyback' biological fitness. One might even consider the possibility that memes would have a 'strategy' of 'parasitizing' biological fitness. Does this prove that memes for agriculture exist, or that agricultural memes exclude the possibility of contributing to biological fitness? Certainly not. The attempt made here is only to demonstrate that memetic algorithm are not necessarily at odds with biological fitness, and that it is at least possible to look at memes in this way.

While agriculture is considered a co-evolutionary phenomena between ants and their cultivars, this is generally not the case with human agriculture. From an evolutionary perspective, the fungus an ant species cultivates could be said to exploiting some behavioral aspect of the ant system, just as much as the ants are exploiting the fungus.

Consider this: Jonathan Swift's infamous traveling character, an astonished Gulliver, has met the Houyhnhnms - a race of intelligent horses who have language and keep domesticated primitive like humans called Yahoos. Gulliver regales his Houyhnhnm host with descriptions of horses, a species that humans (i.e. Yahoos) have domesticated and dominated back in Gulliver's world. To 'quote' Gulliver: " I told him we had great numbers [of horses]; that that in summer they grazed in the fields, and in winter they were kept in houses with hay and oats, where Yahoo servants were employed to rub their skins smooth, comb their manes, pick their feet, serve them with food, and make their beds". The Houyhnhnm host wisely replies: "I understand you well...it is now very plain from all you have spoken, that whatever share of reason the Yahoos pretend to, the Houyhnhnms are your masters. I heartily wish our Yahoos would be so tractable".

Saying that humans are the sole beneficiaries of exploiting certain species of plants and animals is like saying that certain species of plants are exploiting humans, and have increased their fitness by tricking humans into dispersing their genes for them, or that some animals have taken advantage of people by getting them to feed and protect them. Neither is completely true and thus, we maintain that agriculture is a process of coevolution between species, and unlike ants, involves memetic replicators as well as genetic ones in human agriculture.

So how can a forager system become an agropastoral system? One useful scenario may be to look at a hypothetical scenario where foragers adopt a new meme or set of memes - such as management of domesticates. This is where the idea of replicator systems can have great utility for providing explanations of co-evolutionary relationships. A foraging system consists of people collecting resources from the 'wild', or the environment they live in mainly for subsistence purposes. It is generally a stable system in terms of subsistence, because it is generalized or 'unspecialized', and therefore relative yield from their foraging activities overall are more or less stable. On the other hand, agriculture is a highly specialized system which increases variance in the relative yield of resources. How can one element of one system be incorporated by another? Incorporation of a new element into existing systems from this perspective could explain the often piecemeal appearance of domesticate adoption seen in the archaeological record across much of the world.

First, issues of the cost and benefit to replicating the meme should be addressed in terms of production and consumption. It is one thing to obtain or 'consume' a product (an element of a system), and quite another to produce an element (replicating the algorithm). One hypothetical approach would be the proposition that the phenotypic element is consumed (a sort of memetic viral strategy) until the benefits of local production outweigh the costs of distant consumption, allowing/requiring the meme to be transmitted. There's a cost to meme replication if indeed it's an algorithm for information to produce something. People have to take the time to learn the algorithm. Also memes tend to create systems of organized units, or are atomized meaning that information is specialized and compartmentalized for the production or expression of a particular element. Maybe it's too expensive to always immediately replicate a meme from individual to individual, and also, reverse engineering a meme algorithm can be too costly. The cost of meme selection will vary from meme to meme (perhaps this can be explained in terms of an r vs. k selection strategy). But let's say hypothetically there is not as much cost to consume or acquire the element as replicating/transmitting the meme. The element itself can be part of a memetic reproductive strategy for increasing the odds of algorithm replication (especially considering the "less digital" nature of meme algorithms). This is related to the idea of memes as cultural viruses (after Cullen). If the product is

successful enough (ie consumption of it increases) the odds are higher that meme will become replicated, especially when the costs of consumption of it outweigh the cost of local production. When this threshold is reached, someone bothers to learn the meme algorithm for him or herself.

Secondly, adoption of domesticates from which over time might be expected to result in evolutionary change.

This is where the concept of Romer's Rule is useful to consider. Romer's Rule states that an amphibian is just a fish desperately trying to remain a fish, and is meant to describe the conservative nature of evolutionary change. In our example here, a 'farmer' might just be a forager trying to remain a forager. This statement of course implies active and intentional resistance of adopting domesticates by foragers, and in reality no such intention is implied, though it does seem possible. What it does suggest, is that foragers will incorporate a new element into their system without any conscious intent of changing it, though change is inevitable through intensification of certain subsistence behaviors.

A useful example might be the scenario of the adoption of domesticates by Northern European Mesolithic foragers. About 6,000 years ago in Denmark, foragers adopted domesticates from their farming neighbors to the south. The principle domesticate involved were cattle and other livestock, while the situation for the adoption of cereal agriculture is less clear and unconvincing. We accept the theory concerning this example, that these foragers turning farmers were foddering the animals through the winter with leaf hay.

Exhibit A: There is significant evidence that indicates a large amount of "wild resource husbandry" on the part of Mesolithic foragers (Zvelebil, 1994). The evidence suggests coppicing and perhaps pollarding of trees, particularly hazel for increasing nut production and browse for wild animals, keeping forests open to attract prey, and even sized twigs for constructing fish weirs for example. It seems that there was extensive forest management on the part of Danish Mesolithic foragers.

Enter exhibit B: Around 6,000 years ago, there is evidence of domesticated animals, principally cattle and pigs, in what are clearly in every other respect based on ceramic and lithic typologies, Mesolithic sites. Subsequently during a 'transitional' phase, TRB type ceramics and lithics appear. Overall during this transitional period, the material record still indicates a hunter-gatherer lifestyle where wild foods make up the vast majority of the diet. One could consider these new TRB farmers as foragers with a few cows. Production of leaf fodder for livestock at this point would simply be an extension of Mesolithic forest management.

In this example, the new element added to the forager system is livestock. Adding livestock to the system is feasible because there is little cost to adding it to an already stable and successful system. The low cost can be attributed to the compatible element(s) in the forager system of forest management technology (coppicing/pollarding trees). Therefore, the "cow" element does not disrupt the existing system and confers fitness by perhaps providing a more reliable source of animal resources (as well as possible new 'renewable' resources such as milk, wool, etc.), or elevating social status in those individuals in possession of livestock. It seems likely that both benefits are possible.

It's possible that the element(s) related to livestock, were initially 'consumed' before the algorithms for keeping livestock were transmitted. The initial presence of the elements themselves rather than the algorithm in the foraging system may have spurred changes in social views of status and wealth of individuals in relation to others, or 'piggybacked' existing memetic expectations of individual status and wealth which in turn may have affected the degree of competition for status and wealth between individuals. It could be that once a certain level of competition is reached transmission of the algorithm itself from one system to the other took place.