Gatherer, D. (1998). Why the Thought Contagion Metaphor is Retarding the Progress of Memetics. Journal of Memetics - Evolutionary Models of Information Transmission,**2**.
http://cfpm.org/jom-emit/1998/vol2/gatherer_d.html

Why the `Thought Contagion' Metaphor is Retarding the Progress of Memetics

Derek Gatherer
School of Biomolecular Sciences
Liverpool John Moores University
Byrom St.
Liverpool L3 3AF
D.Gatherer@livjm.ac.uk

Abstract
1 - Introduction and Summary
2 - Thesis 1: Memes - To Have and Have Not

2.1 - Dawkins A and Dawkins B
2.2 - Memes in Action
2.3 - Two Difficulties

2.3.1 - Difficulty no.1 - How do we observe a (Dawkins B) meme?
2.3.2 - Difficulty no.2 - How do we quantify a (Dawkins B) meme frequency in an observed population?

2.4 - Dawkins A is Better than Dawkins B
2.5 - Further Difficulties with the Meme-Host Relationship
2.6 - How do we Quantify the Frequency of a Conceptual Meme?

3 - Thesis 2: Lies, Damned Lies and Memetic Transmission

3.1 - Beliefs Concerning a Simple Statement
3.2 - Belief at the Dawn of Memetics

4 - Belief Systems as Complexes of Behaviours
5 - Social Contagion Theory and its Relevance to Memetics
6 - Transmission of Innovations
7 - Who Uses Which Definition?
8 - Criticisms Anticipated

8.1 - That Memes are Only Abstractions
8.2 - That Memetics is Still in its `Mendelian' Stage, and We Will One Day be Able to Identify Memes Directly
8.3 - That All the Above is Inadmissible as it Constitutes an Attack on Memetics
8.4 - That Memeticists Don't Say that People Have Memes Anyway, We Say that Memes Have People
8.5 - That Memes are Only the Homoderivative Subset of Mnemon Replication Events
8.6 - That I Contradict Statements Made In My Earlier Work

9 - And Finally....Yet Another Meme Definition
Acknowledgements
References

Abstract

The most generally accepted definition of the meme, as a `unit of information residing in a brain' (Dawkins 1982), implies a meme-host duality which is the basis of many current developments in memetics, in particular the notion that the passage of such memes (or homoderivative mnemons, following Lynch 1998) from mind to mind constitutes a process that may be considered as `thought contagion'. A critique of religious belief and other non-rational systems of thought, as `mind viruses' (Dawkins 1993), has been built upon such a meme-host duality. This paper provides two objections to the `thought contagion'/`mind virus' theory: a) that the concept of a transmitted belief, as opposed to transmitted information, is highly problematic, and b) that in any case the concept of a meme-host duality is equally suspect. It is suggested that the least philosophically problematic constitution for a science of memetics would be to adopt a behaviourist stance towards memes, to restrict the use of the term to those replicating cultural phenomena which can be directly observed or measured (Benzon 1996). This would release us from the difficulties of the indefinable meme-host relationship, and also have the merit of making memetics more directly comparable to animal behavioural ecology, to the existing branch of social psychology known as social contagion theory, and to the sociological field of empirical diffusion studies.

Key words: meme, belief, religion, linguistic behaviourism, mind virus, thought contagion, social contagion, diffusion of innovations.

---

1 Introduction and Summary

This article is designed to complement the on-going debate concerning the definition of the term `meme', by investigating some of the limitations consequent on certain usages of the word. Broadly, my intention is to show that what is probably the most popular, even one might say the `orthodox', definition of the meme as a `unit of information residing in a brain' (which I shall term Dawkins B, see Dawkins 1982), presents us with serious philosophical difficulties that may hamper (and indeed are already hampering) the development of memetics as a science. I shall argue that an earlier and broader picture of the meme as a `unit of cultural transmission, or a unit of imitation' (which I shall call Dawkins A, see Dawkins 1976), is in many respects a better working definition, despite the fact that it is now generally regarded as obsolete. The conclusion of this argument is that the best current definition is that of Benzon (1996) although, as with any definition, there is scope for refinement.

I shall focus on two main theses, which develop as a consequence of discussing the difficulties that lurk within Dawkins B:

1. There can be no population memetics.
2. Memetics cannot be used to study why beliefs spread.

In brief, thesis 1 is consequent on the impossibility of quantifying meme frequencies in populations,

since this requires a clear meme-host relationship and that is precisely what we cannot obtain in most cases, as will be shown. Thesis 2 is consequent on the impossibility of accurately identifying when belief has been replicated or transmitted. It is possible to accept thesis 2 while still rejecting thesis 1. However, if thesis 1 is accepted, thesis 2 is a necessary consequence.

Some may be horrified at thesis 1. Part of the initial appeal of memetics was that it seemed to suggest that the powerful theoretical tools of population genetics might be applied to revitalise the venerable but rather moribund field of cultural evolution. The present author has also subscribed to that view (Gatherer 1997a, b, c, 1998) so this article is partly a self-criticism. If we can have no *population* memetics, some might feel that there is no point in having *any* memetics. However, memetics may still be scientifically applied in other ways, and may still make a contribution to the humanities. In particular, I shall attempt to argue that memetics may be best constituted as a science by adopting a behaviourist perspective. By behaviourist, I do not mean the rigid `neo-behaviourist' tradition of Skinner (eg. Skinner 1972), but in the broader sense of Watson (see Cohen 1979). Under the terms of such a `soft' behaviourism, only observable behaviours and artefacts would be considered, and memetics would cease to concern itself with unobservable mental entities such as beliefs or thoughts. Benzon (1996) has previously come to this conclusion, but by a different route. This critique is particularly directed against those who see memetics as a science of `thought contagion' (eg. Lynch 1996a), presumably extrapolating from the well-established branch of social psychology known as social contagion theory. This latter dates back to LeBon's `Psychologie des foules' in 1895, and Baldwin's study, in 1897, of what later became known as the `Werther Effect' (apparently contagious suicidal behaviour, see Phillips 1974; reviews of social contagion by Hamilton and Hamilton 1981, Levy and Nail 1993). Importantly though, social psychology is largely (although admittedly not exclusively) constituted as a behaviourist and empirical science, and thus avoids many of the philosophical difficulties that `thought contagion' theory finds itself in.

The first step is to examine why: a) we cannot have a meme-host duality, and why therefore: b) we cannot quantify meme frequencies per capita, and why we therefore: c) cannot have any population memetics.

---

2 Thesis 1: Memes - To Have and Have Not

2.1 Dawkins A and Dawkins B

I shall refer to Dawkins' original formulation of the meme as Dawkins A (Dawkins 1976). Several years later, Dawkins presented a revised version of the meme which I shall refer to as Dawkins B (Dawkins 1982). According to Dawkins (1982), the revised version was intended to clarify the relationship between memes and cultural products. Here it will be argued that the Dawkins B definition did not clarify the situation but trapped memetics in a conceptual impasse from which it has not yet emerged.

For reference, the original formulations are as follows:

> ***Dawkins A****: `..a unit of cultural transmission, or a unit of imitation.' (Dawkins 1976, p.206); `Examples of memes are tunes, ideas, catch-phrases, clothes fashions, ways of making pots or of building arches' (Dawkins 1976, p.206); `Popular songs and stiletto heels are examples. Others, such as the Jewish religious laws...' (Dawkins 1976, p.209); `Perhaps we could regard an organised church, with its architecture, rituals, laws, music, art and written tradition, as a co-adapted stable set of mutually-assisting memes.' (Dawkins 1976, p.212);*

> *`Memes for blind faith have their own ruthless ways of propagating themselves.' (Dawkins 1976, p.213).*
>
> ***Dawkins B**: (referring to the original Dawkins A definition, above) `..I was insufficiently clear about the distinction between the meme itself, as replicator, on the one hand, and its `phenotypic effects' or `meme products' on the other. A meme should be regarded as a unit of information residing in a brain..... It has a definite structure, realized in whatever physical medium the brain uses for storing information....I would want to regard it as physically residing in the brain.' (Dawkins 1982, p.109); `The phenotypic effects of a meme may be in the form of words, music, visual images, styles of clothes, facial or hand gestures.....' (Dawkins 1982, p.109).*

The central flaw in the `thought contagion' and mind virus' hypotheses may be summed up in a single phrase: these theories require individuals to *have* memes. In Dawkins B, the state of having a meme is taken as synonymous with the state of having a certain informational pattern in the brain. But, as will be demonstrated, this is a poor basis for building any theory of cultural evolution. Others have criticised memetics by claiming variously that the meme is a useless abstraction, or that memetics has nothing to contribute to the theory of cultural evolution, or that culture does not evolve (eg. Rose et al. 1984, Barbrook 1996, Gould 1996, Harms 1996, Pinker 1998). I do not support such criticisms, which call the validity of memetics into question. *My* criticisms, by contrast, are intended to strengthen the theoretical basis of the memetics movement. If we are to have a cultural evolutionary theory, then we need units of replication/selection, and the meme concept would seem to be as good a way as any of approaching this problem. The mistake lies in the frequent assumption that individuals *have* memes. But because, as will be shown, we cannot assign memes *to* individuals with sufficient reliability or regularity, we cannot produce meme frequencies, defined as the proportion of individuals in a population that possess or have a meme. Gene frequencies *are* absolutely necessary to population genetics, and if there is to be a population memetics we correspondingly need to have unambiguous meme frequencies. There are some circumstances in which we can derive a statistic of this sort but, as will be demonstrated, this kind of `meme frequency' statistic, if we are to have it at all, must be defined in a purely behaviourist manner.

By contrast, individuals do *have* genes. They may or may not pass them on to the next generation of individuals. However, whatever the reproductive success or otherwise of an individual, that individual carries those genes around in its body for the duration of its life. A population geneticist may derive allele frequencies as the proportion of individuals who carry one gene sequence or another. Gene frequencies are of course abstractions, pure quantitative values, but they are abstractions which relate directly to a physical reality, ie. the reality of gene sequences within the bodies of individuals. Thus, providing the required technical methods are available, one may derive an unambiguous estimation of gene frequencies which are comparable between one population and another.

Memes, on the other hand, are more difficult to pin down. Dawkins A included concepts, cultural artefacts and subjective states such as beliefs, whereas Dawkins B restricted memes to units of neural information giving rise to behaviours or the production of artefacts (Dawkins 1982, p.109). Dawkins is careful to stipulate that when a meme is transmitted from one brain to another, it is not necessary that exact neural configurations are reproduced. Two individuals who are exhibiting the same mental state, entertaining the same idea, performing the same behaviour, are taken to have the same meme, for all intents and purposes - even if their neural configurations are not identical. Thus at this most basic physical level, Dawkins B memes are less tangible than genes. They are merely informational specifications within brains.

2.2 Memes in Action

Let us take an example to see how such a Dawkins B framework would apply to a simple system. There is a method of tying a necktie which is called the Windsor knot (named after King Edward VIII). Windsor knots provide a better knot volume and more elegant shape than the standard method of tying a necktie. There is even a mutant form, the half-Windsor, which is slightly simpler and which many people tie in preference to the full Windsor. If we first apply Dawkins A, the knot itself is a cultural replicator, and therefore a meme, as is *also* the idea of tying the knot. *Both* the Windsor knots themselves and the neural activity involved in tying them are taken to be capable of spreading through populations (for instance, the Windsor knot has made a distinct comeback now that ties are a little wider again). By contrast, Dawkins B would restrict the meme to the neural correlates of knot-tying activity, with the knot itself simply becoming the cultural artefact, or `meme product' (Dawkins 1982, p.109) consequent on the activity of the meme in the mind/brain of the tie-wearer.

To reiterate, in Dawkins B, two people who know how to tie a Windsor knot have the same meme, even if their neural configurations producing this activity are widely different. Likewise, when a Windsor knot is taught, the pupil need not have the same neural configuration as the teacher. The important thing is that both individuals can think of the Windsor knot, and can tie it correctly. Under these circumstances, it is said that meme transmission or replication has occurred.

2.3 Two Difficulties

2.3.1 Difficulty no.1 - How do we observe a (Dawkins B) meme?

Two difficulties arise as a consequence of this Dawkins B definition. The first of these is the less serious and is acknowledged by Dawkins (1982), although he gives little indication that he thinks it *is* a difficulty. This difficulty is that the flexibility that he allows in neural configuration corresponding to a Dawkins B meme, the fact that it is *just* a unit of information, presents us with no means of identifying memes separate from their manifestation in the outside world, ie. separate from their `meme products'.

Even if one were able to use some method of scanning an individual's brain to discover the internal neural configuration, we would not be able to identify the Dawkins B memes of that individual. As far as neurobiology is currently concerned, Dawkins B memes must be abstractions. This is not a problem as far as mathematical modelling is concerned, nor is it a problem for systems theory-oriented memeticists (eg. de Winter 1984). However, for those with a more empirical bent, those who would like to study memes `in the wild', it is evident that memes, *even under the rules of Dawkins B*, are only defined *by their manifestations*. Those of us who wish to do empirical memetics are thus left in the position of looking at `meme products', which *are* admissible as actual memes under the more flexible rules of Dawkins A. Little therefore seems to be achieved by the `refinement' of Dawkins A to Dawkins B, other than to require that we continually remind ourselves that the cultural artefacts we are studying are not the `real' memes but that these `real memes' are unobservable informational entities in some brains somewhere.

This first difficulty, although inconvenient, is not fatal. It produces a certain philosophical awkwardness, but the general schema of Dawkins B would still be functional if we were able to assign the abstract memes unambiguously to individuals, and thus derive meme frequencies per capita of human population. But this is something we *cannot* do, for the following reasons.

2.3.2 Difficulty no.2 - How do we quantify a (Dawkins B) meme frequency in an observed population?

This second difficulty is considerably more serious, since it disqualifies both the population memetics and `thought contagion' models of cultural change. Both of these *depend* on a meme/carrier-of-meme or meme/host relationship, a necessary consequence of Dawkins B. In the population memetics approach, the meme is regarded as coding for a meme product which may or may not contribute to the fitness of the individual that carries it. In the thought contagion model, the relationship is considered as more that of virus/host. This is more than just a nicety of semantics, since the decision to adopt one or the other basic perspective can lead to different abstract mathematisations of memetics. For instance, compare Cavalli-Sforza and Feldman (1981) for a population memetics approach where memes (or `cultural traits' as they term them) are primarily transmitted `vertically' within families in a manner precisely analogous to the vertical transmission of genes and an iterative generational model is used with an emphasis on fitness and selective advantage; versus Laland (1992), Lynch (1998) and Takahasi (1998) for contagion approaches which emphasise `horizontal transmission' and the reproductive, rather than selective, efficacy of the meme. Cavalli-Sforza and Feldman (1981) also present models of the contagion variety, where transmission of the `cultural trait' is horizontal between siblings and unrelated individuals, and the mathematics uses equations from epidemiology rather than population genetics. As Wilkins (1998) points out, both the epidemiological and population genetics methods are derived from the same basic mathematical models of differential reproduction, but the emphases in each case are different.

However, whatever the relative merits of the mathematical systems, in neither case is it possible *in practice* to accurately derive any quantification of meme-host relationships. The example of the Windsor knot will again serve as an illustration. Let us imagine a room of 100 individuals, only one of whom can tie the Windsor knot, and that individual is the only one wearing the Windsor-knotted tie. If we consider the frequency of the Windsor knot meme to be the frequency of the actual occurrence of the knot (taking a Dawkins A-type interpretation), then we have a frequency of p=0.01. If, on the other hand, we take a more orthodox Dawkins B interpretation, and consider the frequency of the Windsor knot meme to be the frequency of the neurally encoded ability to tie the knot, ie. the mental informational unit corresponding to the knot, then again we have a frequency of p=0.01. The two approaches would seem initially to be in agreement.

This situation becomes rapidly more ambiguous when our tie-wearing individual proceeds to tie the knot around the necks of all the other individuals in the room. The frequency of our meme as defined by the occurrence of the knot (Dawkins A) becomes p=1. However, the frequency of the neurally encoded ability (Dawkins B) remains p=0.01 (unless we allow that some of the individuals learned how to tie the knot while the process was being performed on them, in which case it will increase).

This thought experiment demonstrates that Dawkins A has the advantage of telling us about the change in frequency of a cultural artefact. However, it tells us nothing about any neural events underlying that. We have a change in Dawkins A meme frequency from p=0.01 to p=1, but that does not reveal the fact that our knot-tying individual is the only one with the knowledge of how to tie the knot. Conversely, under a Dawkins B interpretation, we have no change in meme frequency, but a radical change in the cultural situation of our population. Thus our choice of meme definition places us in two very different situations regarding analysis of the above thought experiment. Furthermore, under the Dawkins B schema we cannot *in any case* identify memes objectively, but only *infer* them from their `meme products'. If we do not see the single individual tying all the knots, we might incorrectly infer that the other 99 individuals in the room have acquired the Dawkins B meme for knot-tying, and mistakenly place our Dawkins B meme frequency value at p=1. Again, if one of the other individuals in the room secretly knows how to tie the Windsor knot but is not disposed to divulge this fact, then we actually have a Dawkins B meme frequency of p=0.02, but we cannot ever know this empirically. Only a selective event (eg. a million dollar prize for Windsor knot tying ability) might suddenly reveal the real Dawkins B meme frequency.

2.4 Dawkins A is Better than Dawkins B

To recap, Dawkins A provides us with information about the *cultural state* of our group of 100 individuals, but tells us nothing about what is going on inside their heads. Dawkins B tells us nothing about the change in cultural state of our group, but rather focuses on *unobservable and merely inferred events* going on inside their heads. Since memetics is a theory of cultural evolution, Dawkins A is preferable as *it allows us to look at culture*. Dawkins B diverts us into conjectures about which memes reside in whose heads, regardless of the objective cultural state of our population. This is the principal reason why there has been little or no empirical memetics in the last 22 years. Dawkins B is not a sound basis for any empirical science, as it relegates the observable, objective manifestations of culture to a subsidiary status. It takes all the data we can gather and dismisses it as mere `meme product', and leaves us tangled in endless speculation about unobservable mental units of information.

Dawkins A, by treating the objective bits and pieces of culture as memes, frees us from the requirement to chase unobservable entities. It also frees us from the meme-host relationship, since Dawkins A memes, as artefacts and behaviours, need not have hosts. Indeed artefacts definitely do not have hosts. Of course Dawkins A is a broad definition and *also* classifies unobservable mental entities as memes. But Dawkins B turns the unobservable mental entities into the *main objects of study*.

2.5 Further Difficulties with the Meme-Host Relationship

Recall the thought experiment involving the room of 100 individuals and the Windsor knot. We saw that in order to identify accurately the true number of individuals who can tie the Windsor knot, ie. the frequency of the hosts of the Windsor-knot-tying meme under the rules of Dawkins B, we might have to produce selective conditions which reveal the full extent of ability to tie the knot. However, even if we can do this, we still cannot say that the individuals capable of tying the knot are *permanent* hosts for any Windsor knot Dawkins B meme, in the way that they may be hosts for any particular gene. All that we *see* is knot tying activity, and a resulting change in the frequency of knots. But knot-tying activity alone (and what other evidence do we have that the individual `has' the Dawkins B meme?) does not give grounds for postulating any permanent meme-host relationship.

Hewlett and Cavalli-Sforza (1986), in their extensive study of cultural traits in the Aka pygmies, adopt just such a technique of asking who can perform a particular skill (50 skills are considered ranging from hunting techniques, fruit gathering, and sexual behaviour, to songs and dancing). They are therefore able to build up a picture of who is able to do what, and from whom they learned it. The result, however, tells us more about social structure and relations than it does about selection pressures or any changes in the cultural state of the society under consideration (except for the observation that the Aka have recently abandoned hunting with bow-and-arrows for hunting with the more lethal crossbow - but note that in this single instance it is no longer ability but an artefact that is under consideration). Notably, many of the skills are possessed by high percentages of the population. For instance 99% of adult males know all net-hunting techniques, 98% know all food gathering techniques, 100% know all sexual techniques. Among adult females, percentages are equally high. The only distinct partition among the adult population is that females are less knowledgeable concerning hunting (only 59% know all hunting techniques) while males are less knowledgeable about food preparation (only 47% of adult males are any good in the kitchen as compared to 100% of adult females). The homogeneity that arises from Hewlett and Cavalli-Sforza's analysis is precisely a consequence of their decision to examine ability to carry out skilled activities, as opposed to a more behaviourist approach of observing and quantifying performance of skills over time. Referring back to the thought experiment above, they are in the same position as one who attempts to quantify knot-tying *ability*.

A later paper from Cavalli-Sforza's school (Guglielmino et al. 1995) adopts a wider perspective and compares the total cultural pools of different populations. This approach does not require the difficult task of assigning Dawkins B memes (or cultural traits) *to* individuals, since a population is deemed to possess a cultural trait if that trait, usually defined as an artefact or behaviour, exists within it somewhere.

As we shall see later, to postulate that anyone who *can* perform a certain cultural activity, *has the meme for* that activity, leads us into even deeper problems where conceptual memes are considered.

Under Dawkins B, there is no easy way to circumvent this problem. Among the unsatisfactory alternatives, one might postulate *another* meme for `wearing the Windsor knot', independent of the one for tying it. This might enable us to say that the frequency of the Dawkins B `knot-tying' meme in our room is p=0.01, and remains so, but the frequency of the Dawkins B `knot-wearing' meme goes from p=0.01 to p=1. But what is the neural process that underlies merely passively wearing the knot? Whereas one might easily concede that performance of a skill involves a specific neural activity, which may be reproduced every time the skill is performed, passive states are less easily coupled to neural processes. A similarly difficult example is the cultural trait for `acceptance of the Fascist salute' (Cavalli-Sforza and Feldman 1981, p.64). *Tying* the knot, or *giving* the salute, are behaviours which presumably must be underlain by neural activities of some sort, unobservable though they certainly are. But what neural activity underlies simply wearing a certain tie, or accepting a certain salute, or travelling to work by train? One may of course easily quantify numbers of individuals travelling by train, or the number of individuals wearing ties, but these statistics would *not* represent Dawkins B memes. Thus a `tie-wearing' Dawkins B meme is not a solution to the problem.

2.6 How do we Quantify the Frequency of a Conceptual Meme?

With conceptual memes, the problems become even more difficult. The memetics of concepts, specifically religious concepts, is considered by Dawkins (1993), but the major venture into this territory is by Lynch (1998). Lynch uses the term `mnemon' to indicate an abstract unit of memory.

> *`Thus, the principle abstractions manipulated with memetics theory are memory abstractions, or mnemons' (Lynch 1998, Section 4).*

This is not dissimilar to the Dawkins B meme, the `unit of information residing in a brain', although of course one might insist that Lynch's mnemons refer only to the memory aspect of neural architecture. What is clear is that Lynch considers *awareness of* simple statements, or *belief in* simple statements, to be mnemons, and that the transmission or replication of such mnemons is the basis of the process of `Thought Contagion' (Lynch 1996a).

Thus Lynch makes statements such as the following Lynch 1998, Section 12):

> *`Take the mnemons expressed by the following 3 statements, for instance:*
>
> *Mnemon A: "There is only one true God".*
>
> *Mnemon B: "Christ is Lord".*
>
> *Mnemon C: "Unbelievers are damned".*
>
> *Mnemon D: "Earthly life is better among believers".'*

[There are of course actually four of them, but I quote verbatim. As I have attempted to demonstrate

throughout this article, `Thought Contagion' theory pays scant regard to the accurate quantification of memes].

These statements, or more exactly *awareness of* these statements, are taken by Lynch to be encoded or somehow instantiated in memory. The reader is referred to his article (Lynch 1998) for further examples involving birth control, church attendance etc. Individual concepts are thus taken as mnemons which exist in the heads of their hosts.

Let us accept this system at face value for the moment, and consider such a concept-mnemon: `Napoleon died in 1821'. This has been physically replicated in print many times throughout history books. An individual who is aware of this might be said to `have' this meme/mnemon, under Dawkins B or Lynch. However, this individual is also capable of transmitting a vast variety of other related memes/mnemons, such as `Napoleon died in 1820', `Napoleon died in 1819' and so on. Now does this individual *have* these mnemons as well? The fact that the first option is factually correct and the other two are false merely perhaps makes it more likely that the first option will be transmitted (perhaps in the selective context of a history exam). However, any individual who has heard of Napoleon *at all* will be *capable* of constructing and transmitting an infinite variety of mnemons of the form `Napoleon died in *x*'. If we accept that any individual who is *capable of tying* a Windsor knot *has* the Windsor knot meme (and the Dawkins B interpretation *would* seem to imply this, and so would Hewlett and Cavalli-Sforza 1986), then we must accept that anyone who is *capable of transmitting the concept* `Napoleon died in 1819', *has* that mnemon: `awareness of the concept that Napoleon died in 1819'. If we can transmit something, we must be aware of it, unless we are transmitting in an unconscious state.

Similarly, to return to Lynch's example, anyone who has the mnemon A, above: `There is only one true God', must also have the mnemons: `There are only two true gods', `There are no true gods' etc, since that individual would be capable of transmitting these mnemons. Even if that individual does not actually transmit any of the mutant derivatives of the first mnemon, that person would still be *aware of* the alternatives and thus, according to Lynch's rules, must *have* the mnemons for those alternatives. Since the alternatives are in principle infinite (ie. `there are only *x* true gods', for *x*= 0 to infinity), Lynch's system of conceptual mnemons, taken as *awareness of* any particular statement, inevitably leads us into a picture of the human mind as host to an *infinite* array of potentially transmitted concepts. This is clearly unsatisfactory.

The problem of dealing with conceptual mnemons/memes as stored statements, or stored awareness of statements, is a further strong stimulus to abandon the strained host-meme correspondence. We have to consider memes as cultural entities which do not *belong to*, or parasitise, any individual. For this reason, Dawkins A is again a preferable definition. Individuals may be involved in the *production* of memes, but those memes/mnemons are *not* resting latent in the individual as a gene does. This position will be developed further once we have considered thesis 2.

In summary, since memes (defined by Dawkins B) are not directly observable, we can only study their meme products. Furthermore we cannot assign Dawkins B memes to individuals with any reliability since the relationship of a meme product, whether behaviour, artefact, or concept, to the hypothesised unobservable Dawkins B meme that produces it, is always a difficult one. Consequently, we cannot derive meme frequencies per unit of human population with any reliability. Without such reliability, there can be no population memetics.

---

3 Thesis 2: Lies, Damned Lies and Memetic

Transmission

3.1 Beliefs Concerning a Simple Statement

Thesis 2 is that memetics cannot be used to study why beliefs spread. Note that the phrase `How Belief Spreads Through Society' is the subtitle of the book on memetics by Lynch (1996a).

One possible difference that might be posited between the mnemon `Napoleon died in 1821' (N1821) and other mnemons of the type `Napoleon died in x', is that N1821 *is believed* and the others are not. Perhaps I *have* the meme N1821 *because I believe it*. The study of belief states has fascinated, perhaps obsessed, memeticists ever since Dawkins, and indeed memetics is best know to the general public for the aggressive critique it has provided of religious belief as a `mind virus' (Dawkins 1993).

However, there are also serious problems with the attribution of memetic status to belief states. Lynch (1998) hold beliefs in statements to be mnemons *in their own right*. We therefore have `awareness of the statement that abortion is a mortal sin' as *one* mnemon and `belief that abortion is a mortal sin' as a *different* mnemon. Similarly (Lynch 1998, Section 9):

> *`Mnemon P is the belief that "bee pollen invigorates".*
>
> *Mnemon Q is awareness of the "bee pollen invigorates" proposition.'*

Bearing this in mind, let us further analyse the example of the statement `Napoleon died in 1820' (N1820). If a teacher deliberately transmits this false statement to a pupil and the pupil believes it, then we cannot say that *belief* has been replicated or transmitted, since the teacher does not believe in this statement. According to Lynch's system, the mnemon `awareness of N1820' *has* been replicated, but the mnemon `belief in N1820' has not. The same applies to misunderstandings, where the teacher attempts to transmit the correct statement (N1821) but the student ends up believing an incorrect version. Transmission of belief is clearly far more problematic than transmission of information. Where transfer of information alone is concerned and belief is disregarded, the misunderstanding is merely a mutational event, an error of information flow. Likewise, deliberate lying would still be an accurate transmission of information, regardless of the credence given to that information. Despite these difficulties in the transmission of belief, memeticists frequently give the impression that transmission of belief is a straightforward matter of the replication of one person's belief in the mind of another (eg. Dawkins 1993, Lynch 1996a).

No matter how often a teacher repeats a statement, there can be no guarantee that the pupil will believe the statement. Many factors may influence the outcome, but that outcome is not a matter of simple replication of the teacher's beliefs in the mind of the pupil. Belief is not transmissible, but is something that happens *after* the transmission of information. While we *can* often observe the communication of information, we can *never* directly observe transmission of belief. *Information and belief are not the same kind of thing.* Lynch's system, that has `awareness of x' *and* `belief in x' as two variant forms of *the same kind of thing*, ie. as mnemons, is based on a highly doubtful ontology.

3.2 Belief at the Dawn of Memetics

Lynch's conflation of information with belief was not the first incidence of this error in memetics. One of the examples that Dawkins (1976) gives of a meme is `belief in life after death', although he actually quotes this from a personal communication by Humphrey (Dawkins 1976, p.206-207). When a meme

`parasitises a brain' (quoting Humphrey again - and here we have one of the earliest uses of the meme-host perspective), it is the transmission of belief that is taken to matter. Religious leaders acquire disciples, then larger sets of followers and eventually a whole church by transmitting the necessary set of memes for their particular religion, secured by other memes such as `faith' and `fear of divine retribution' which are co-adapted to ensure that the other memes are not easily cleansed from their hosts. This is the standard memetic view of religion, initiated by Dawkins (1976) and developed further by Dawkins 1993 and Lynch 1996a. What is under consideration in these analyses is the spread of belief.

Or is it? Dawkins' original discussion soon veers away from belief and into more abstract territory, referring to the `idea of God' Dawkins 1976, p.207). This is not the same thing as belief, in fact is rather akin to the idea of a flat earth, but Dawkins fails to point out the distinction. Like Lynch, Dawkins appears to see awareness of the concept of God, and belief in the actual existence of God, as similar kinds of thing, just two memes - different memes admittedly, but nevertheless the fact that one is a concept and the other a belief is not seen as a serious problem. Later, in discussing scientific theories, he slides from memes-as-beliefs to meme-as-ideas with a disconcerting rapidity. For instance, we soon have: `....all biologists [who]...believe in Darwin's theory', equated with: `..every individual who understands the theory.....' (Dawkins 1976, p.210). The absurdity of this is illustrated by the fact that I understand everything about fundamentalist Christian theology; raptures, pre-millenarianism, dispensationalism, total depravity, predestination and irresistible grace (yes, I am Scottish, in case you were wondering). However I don't believe a single word of it. By contrast, when it comes to Darwinism, I both understand and believe.

To recap on thesis 2, that memetics cannot be used to study how beliefs spread: an individual may have a set of beliefs, but these cannot be memes, since a) they cannot be transmitted. All that can be transmitted is information. Belief is not itself information, but an attitude towards information. By having statements and belief in statements as mnemons, Lynch is saying that the attitude towards a mnemon is itself a mnemon. Additionally, b) the only empirical evidence we have of belief is through behaviour, whether linguistic behaviour (eg. uttering `I believe' - a performative utterance in the jargon of linguistic philosophy), or any other cultural manifestation of belief. One may copy the behaviour, but whether this actually leads to belief is another matter entirely. Therefore beliefs themselves do not qualify as memes/mnemons in any meaningful sense.

---

4 Belief Systems as Complexes of Behaviours

Having stated my two theses: 1, that there can be no population memetics using Dawkins B or the mnemon system (Lynch 1998), since there can be no reliable calculation of meme or mnemon frequencies per unit of population, and 2, that memetics cannot be used to study how beliefs spread, since belief is not transmissible but occurs, or fails to occur, independently in the recipient after transmission of information, I shall attempt to put forward an alternative. This is a behaviourist scheme, which treats memes as cultural events, behaviours or artefacts which may be transmitted or copied. Outside the occurrence of the event, the practice of the behaviour, or the lifetime of the artefact, the meme has no existence. The meme does not `go anywhere' when it is not manifested. It is *not* stored in some neural data bank, some internal meme repository. Lynch (1998, Section 4) would have us believe that the unexpressed meme: `resides very redundantly in someone's brain', but again this conjures up images of vast memory banks of mnemons encoding `awareness of statement x', not a view of the architecture of the brain that is very consonant with current thinking in either neurobiology or linguistics. By contrast, the behaviourist meme *is* its own manifestation. This definition is basically Dawkins A, but with the mentalistic component of that meme definition stripped out. It is a similar definition to that of Benzon

(1996).

Let us take the example of a laboratory scientist, an evolutionary ecologist, carrying out a complex set of daily behaviours: fruit flies may have to be fed, progeny counted, genomic DNA prepared, agarose gels loaded, data processed on a computer etc. The belief system behind this may be Darwinism, and the particular set of behaviours may be designed to quantify genetic differences in fly genomes from different habitats. None of the behaviours in the laboratory make much sense without the presupposition of belief in Darwinian theories of evolutionary change. The scientists also say certain typical things to each other; they count allele frequencies, they discuss selective pressures, they debate the significance of results for the theory. These are transmitted behaviours; anyone hoping to become a successful evolutionary fly ecologist will have to copy and master them all. They are memes in their own right. Likewise, a priest of a religion may perform a standard set of ritual activities, say standard things in a religious context, live in a certain way and encourage her flock to do likewise (or perhaps to have lots of children). Novice priests have a lot of memes to copy too. These things presuppose the belief system. Religions, political doctrines and scientific theories may be referred to in terms of the behaviour they produce. This is not quite the same thing as saying that they are defined in terms of this behaviour, but where memetic transmission is concerned, much of what is transmitted is behavioural. In both cases, however, there is still room for individuals who perform the behaviours without believing the underlying theory. There are sceptics and jokers, subversives and cynics, in every profession.

Nevertheless, behaviour, and the artefacts produced by behaviour, such as micropipettes, laboratories and cathedrals, *are* more easily quantified than conceptual abstractions. It is possible to measure the increase in the number of Christian churches per capita in the Roman Empire, to measure the number of scientific articles on the subject of Darwinism per total articles published, to measure the percentage of Aka hunters carrying crossbows, to measure the percentage of votes cast for political parties defining themselves as socialist etc. *These* are our meme frequency statistics, or as near as we shall ever get to such a thing. *These* memes are behaviours, or artefacts that are the products of behaviour, and *not* abstract informational instantiations in individual brains.

And, crucially, individuals do not *have* any of these memes. They build them, say them, do them, make them, assent to them or deny them, but the memes are entirely outside the human beings that generate them. These meme frequency statistics are *not* per capita of human populations, and therefore *do not constitute a body of data which is formally analogous to that of population genetics*. For this reason, there can be no population memetics.

---

5 Social Contagion Theory and its Relevance to Memetics

In this section I ought to be clear that I am discussing the social contagion school of social psychology, and not the thought contagion wing of the memetics movement.

Social contagion psychologists refer to a wide variety of socially transmitted phenomena. These can include everything from major social problems such as aggression (Wheeler and Caggiula 1966), mass hysteria phenomena (Stahl and Lebedun 1974) and antisocial behaviour (Sigelman and Sigelman 1976), right down to such things as contagious laughter (Freedman and Perlick 1979), jaywalking (Lefkowitz et al. 1955), hyperventilation (Moss and McEvedy 1966), binge eating (Crandall 1988) and coughing (Pennebaker 1980). The reader is referred to Levy and Nail (1993) for further details and a panoramic

review of the field.

One theme of contagion research is the modification of epidemiological modelling to the particular phenomenon under consideration (Hamilton and Hamilton 1981). In most cases however, the approach is purely empirical and no mathematization is presented. Additionally, social psychologists tend to have little to say about evolution. Social contagion phenomena are either generally seen as unimportant to the evolution of society as a whole or, perhaps more frequently, the social psychologist is raised in a tradition which rejects the entire notion of social evolution. Thus in `traditional' sociology and social psychological contagion theory, there is little if any reference to selective pressures, competition etc.

Social contagion theorists have tended to split their field into three areas: disinhibitory, echo and hysterical (Levy and Nail 1993). In disinhibitory contagion, what seems to be involved is not the transmission of any meme so much as contagious release of inhibitions which then serves to permit an aspect of the unconscious to manifest itself. The connections with Freudian psychology are obvious (Freud 1959).

Social contagion theory tends to direct itself towards the immediate, spontaneous and transient, such as crazes, manias, and fads. Much of the discourse of pop memetics is concerned with such things. However, if memetics is to explain cultural evolution (and that is surely its aim) then it needs a theory that can deal with more permanent changes in behaviour. Blumer in the late 1930s (cited by Levy and Nail 1993) attempted to enlarge the domain of social contagion theory to include more meaningful social changes such as financial panics, patriotic hysteria, the dancing mania of the Middle Ages and various witchhunts of more recent times. These `social epidemics', as Blumer termed them, can result in more serious consequences, since patriotic hysteria can cause genocide or war, financial panics can wreck economies etc.

It should also be remembered that, for the social psychologist, contagion theory is just one of a variety of ways of explaining the behavioural influence of one individual on another in society. To quote Levy and Nail (1993), social psychologists also have to consider:

> *`...conformity, obedience, persuasion, compliance, deindividuation, social norms, contagion, interdependence, leadership, reactance, social facilitation, social inhibition, social loafing, and vicarious learning...',*

as well as the more obvious factors such as education and indoctrination. Memeticists tend to lump all these factors together into transmission or replication (there may be a subtle difference between these last two but in memetics they are often used interchangeably). However, contagion only represents a small proportion of all memetic events. To claim, as Lynch (1996a, b, 1998) does, that an epidemiology-derived modelling system can be applied to the spread of beliefs through society, is very likely to fall foul of social psychologists, who recognise a more subtle taxonomy of human discourse.

Where social contagion theory enters the evolutionary camp, in the work of Cavalli-Sforza and Feldman (1981), the point is made that innovations tend to spread horizontally whereas long-standing elements of culture tend to be transmitted vertically from elder members of families (and are thus analysed using a population memetics approach). Innovation, especially technological innovation, can be an important fact producing major changes in society (eg. the introduction of motor cars, antibiotics, computers etc.), but once established as part of a culture, the horizontal, social contagion model of transmission is frequently superseded by the vertical, population memetics model.

Crucially, however, Cavalli-Sforza and colleagues almost always tend to study the spread of the empirically observable (Cavalli-Sforza et al. 1982, Cavalli-Sforza 1986, Guglielmino et al. 1995, and to a

lesser extent Hewlett and Cavalli-Sforza 1986). The distinction between traditional social contagion theory (whether evolutionary or sociological) and `thought contagion' theory is that in the latter, mental events are under consideration.

Before leaving the subject of social contagion, consider the phenomenon of mass hysteria, a recent subject of much discussion in the Memetics Mailing List. Participants at last year's most publicised funeral were seen on worldwide television to exhibit waves of weeping, shouting and, much commented on, spontaneous and what one might call `culturally inappropriate' applause. Are we to suppose that the wave of applause represented the passage of mnemons from head to head in a rapid wave of mnemon propagation? And that these mnemons were then immediately translated into the meme product of actual applause? What is happening in a person's head when he hears others applauding and impulsively joins in? What is a `clapping' mnemon anyway? I know what clapping is, I have done it often enough. I do not belong to a culture in which there is no clapping (*then* it might be conceivable to suggest that I have no clapping mnemon). Since I must `have' a clapping mnemon, what mnemon do thought contagionists propose is spreading through a crowd who all already know what clapping is? Is it the `clap now' mnemon perhaps? How does the `clap now' mnemon differ from the simple `clapping' mnemon?

I cannot answer these questions, but if thought contagionists claim that their theory can explain phenomena such as hysterical contagious applause, then they must be able to provide cogent answers to them.

In any case, a more parsimonious approach would surely be simply to adopt a behaviourist stance and just record who is clapping, where, and for how long; no mnemons, no abstractions of memory instantiation, no homogenic instantiation events, *just behaviour*.

---

6 Transmission of Innovations

Another related area which has recently been the subject of much discussion in memetic circles, is that of the theory of diffusion of innovations, principally the work of Rogers and his school. Rogers and Shoemaker (1971) trace its origins back to the work of Tarde in 1903, almost contemporaneous with the origins of social contagion theory.

This field however, is mostly about the production of artefacts and/or the adoption of cultural practices, and therefore cannot be easily reconciled with Dawkins B, or the mnemon approach of Lynch (1998). Typical questions for cultural diffusion studies would be: what is the rate of growth of in the proportion of houses with satellite television and their geographical distribution?, or, what is the growth in the use of catalytic converters and unleaded petrol in motor vehicles?, and so on. Diffusion theory makes no speculation concerning mnemons or other internal elements in human minds.

For example, in the bibliography of 337 classic diffusion theory-related publications, produced by Brown (undated, probably 1964), only four are reconcilable with a Dawkins B-type memetic approach. Only one of these concerns itself directly with religion: Hawley's study of the uneven acceptance of missionary Catholicism in the southwestern USA (Hawley 1946), although two theoretical papers by Rashevsky deal with beliefs, prejudices (Rashevsky 1951) and ideologies (Rashevsky 1952). The fourth paper is a study of the spread of rumours (Rapaport and Rebhun 1952). The last of these subjects is probably equally compatible with a behaviourist approach, since one may simply record who transmits the rumour (verbal behaviour) and who has already been the recipient of an attempted transmission, without the requirement for any mnemons of `having rumour', `believing rumour' etc. Among the more typical subjects for diffusionists are vaccine use, agricultural practices, medical practices, adoption of fluoridation, spread of

cowries as a currency, arithmetical techniques and horse-riding (Brown, undated, probably 1964).

It is significant that neither religion nor belief occur in the index of Rogers and Shoemaker's compendious review of the diffusion research field (Rogers and Shoemaker 1971). These authors list over 1500 articles on diffusion research, but are only able to add one more paper on the diffusion of beliefs, that of Erasmus (1952). Like the social contagion psychologists, diffusionists are overwhelmingly empirical investigators. Their material is the observable and the quantifiable, not the abstract and internal.

---

7 Who Uses Which Definition?

Dawkins B has almost become the orthodox definition of the meme. The principal point of this paper is that Dawkins B represents a retrograde step from Dawkins A, and that it has been responsible for the spawning of all manner of `thought contagion' speculation which is an active hindrance to empirical work in the subject, and a source of ridicule in the scientific community at large (as those of us who are academics can verify).

Those cultural evolutionists who do not describe themselves as memeticists nevertheless often use concepts that are similar to those of Dawkins. For instance, Lumsden and Wilson's definition of the `culturgen' as: `the node of semantic memory' (Lumsden and Wilson 1985, p.348) seems to be very close to Lynch's (1998) mnemon, although importantly, unlike Lynch, they do not construct semantic memory as a stack of `awareness of *x*' modules. Richerson and Boyd (1978) are less specific concerning their `culture-types', which is a unit of selection (not necessarily of replication) in a cultural system:

> *`...natural selection should act on culture-type, increasing the frequency of those items of culturally coded information in a population which.....produce phenotypes that are more successful in passing the culture-type to the next generation' (Richerson and Boyd 1978, p.132)*

If, by `culturally coded information', Richerson and Boyd (1978) imply neural informational units, then this is clearly like Dawkins B. However it is not entirely clear if this is what is implied, since: `..the precise mechanism of inheritance of the code is likely to be largely irrelevant' (Richerson and Boyd 1978, p.132). Whether or not a Dawkins B-type model is implied, Richerson and Boyd do imply a meme-host duality in that they refer to: `cultural codes as they are transmitted from one individual to another' (Richerson and Boyd 1978, p.132), and their mathematical treatment follows from this.

Cavalli-Sforza and Feldman (1981), as discussed above, tend generally towards a Dawkins A definition in theoretical discussion of their `cultural traits' (also see Cavalli-Sforza and Feldman 1983, Feldman and Cavalli-Sforza 1984, Feldman et al. 1985), and in practice tend to be more behaviourist, if somewhat irregularly (Cavalli-Sforza et al. 1982, Hewlett and Cavalli-Sforza 1986, Cavalli-Sforza 1986, Guglielmino et al. 1995).

Those who use more orthodox Dawkins B interpretations or close relatives, include de Winter (1984), Dennett (1990, 1991) and Speel (1997), and it is the source for the computer modelling of a simple society, carried out by Doran (1997). Ball (1984) begins with a resolutely Dawkins B interpretation, but then takes it further in claiming that *all* mental contents are memes, including those aspects of information which are products of Skinnerian conditioning. Memes thus constituted are no longer units of imitation but units of mental content. A similar view is expressed by Gabora (1997), who sees memes as the units of mental experience, down to and including sensory qualia such as `a vivid impression of

red'. Preti and Miotto (1997) also define memes as simply `mental representations'.

Benzon (1996) places memes in the external environment as cultural replicators. Briefly, his central point is that it is artefacts that are copied, and therefore it is artefacts that replicate (albeit somewhat haphazardly). Like Benzon, I do not deny that psychological traits exist - neither of us is a behaviourist in the sense of Skinner or Ryle - but that they are too intangible to be copied in any meaningful sense of the word. The appeal for a `soft' behaviourist memetics here, and in Benzon (1996), constitutes a restriction of the Dawkins A definition to the artefactual, and the elimination of those aspects of Dawkins A which are mentalistic.

Thus reconstituted, human memetics might have more in common with the thriving animal behaviourist wing of the subject. For instance Ficken and Popp (1995), when measuring the evolution of the `gargle vocalisation' of the black-capped chickadee, produce statistics of how many songs of certain types have been produced in a certain period. There is no attempt to assign song memes to individual birds, since individuals have a large repertoire of songs. Similarly, Guglielmino et al. (1995) seek to define which memes (or cultural traits) belong to which cultures, and to work through the evolutionary implications of meme distribution at the societal level. Here we have the meme as artefact/behaviour studied in its own right, without the confusing mentalistic stance that currently plagues human memetics.

---

8 Criticisms Anticipated

Some may reply to the above by saying:

8.1 That Memes are Only Abstractions

One might argue that, since memes are abstractions, we need not worry overly about their ontological status (de Winter 1984). All one needs are tokens for mathematical analysis. This is unsatisfactory, as mathematical modelling is only of value if the model is a reasonable representation of reality. Outside of quantum physics, scientists are generally constrained by the requirement that their models should mirror nature. Even in quantum theory, mathematicians were only driven to the counter-intuitive by the overwhelming weight of experimental evidence. Memetics should be firmly grounded in reality.

8.2 That Memetics is Still in its `Mendelian' Stage, and We Will One Day be Able to Identify Memes Directly

Another similar criticism might be that before the discovery of DNA, geneticists had to be content with a highly abstract definition of the gene. Prior to the molecular era, genes were only identifiable in terms of phenotypes. Thus a fly geneticist could only identify the white (*w*) mutant allele, indeed could only infer its very existence, by the observation of white-eyed flies. Likewise, one might argue that the (Dawkins B) Windsor knot meme is only identifiable by observation of Windsor knots in this present era of `Mendelian' memetics, until we have sufficiently advanced in neurobiology to identify the true, internal, mental, neural configurational, meme, in itself without recourse to the external manifestation - just as we can now identify the white (*w*) mutant allele by DNA sequencing without looking at the flies' eyes at all. This argument fails however, for the following reason.

Recall the thought experiment with 100 individuals and the Windsor knot. The individual with the

knowledge of how to tie the knot was able to produce 100 knots while there was still only one meme (Dawkins B definition) for the Windsor knot. By contrast, in fly genetics, each white eyed fly corresponds to two white (*w*) mutant gene sequences (the allele is autosomal recessive - FlyBase 1998). There is no ambiguity; 100 white-eyed flies equals 200 white (*w*) mutant alleles. But in memetics, we cannot be sure how many Windsor knot memes (Dawkins B definition) were responsible for 100 Windsor knots. Thus the situation in memetics at present is not analogous to the Mendelian era of classic genetics. It is true to say that genes were merely abstractions in those days, but they were a very different kind of abstraction to the meme.

8.3 That All the Above is Inadmissible as it Constitutes an Attack on Memetics

The above two theses, that there can be no population memetics (except in a strictly behaviourist sense) or thought contagion theory, and that memetics cannot be used to study belief states, are bound to cause some consternation within the memetics community. Many may take them as an attack on memetics as a discipline. This is *not* my intention. Rather I see myself as strengthening memetics by preventing its inflation beyond the boundaries of what is acceptable for a scientific field. Population memetics proper (ie. those authors who treat memes as culturally inherited traits) is mostly conducted in a behaviourist manner. Its best empirical results have been achieved when investigators cease attempting to assign memes to individuals and treat them as behaviours occurring within a cultural milieu (Guglielmino et al. 1995).

8.4 That Memeticists Don't Say that People Have Memes Anyway, We Say that Memes Have People

Lynch (1996b) inverts the meme-host relation to give memes which have people, rather than the converse.

> *`In the area of population psychology and psychohistory, memetics achieves just such a paradigm shift by inverting an everyday question. Instead of asking how people acquire ideas, the new paradigm asks how ideas acquire people.'*

This, however, does nothing to help. When does the meme `Napoleon died in 1822' have me? Only when I am transmitting it? Or all the time, given that I transmitted it on one occasion and may or may not transmit it in the future?

8.5 That Memes are Only the Homoderivative Subset of Mnemon Replication Events

Although Lynch (1998) has mnemons of the type `awareness of *x*' or `belief in *x*', not all of these mnemons are memes in his system. Only homoderivative mnemons are memes. By homoderivative, Lynch means that they are produced by replication of the same mnemon in another individual. In Lynch's own words, it is: `A memory item, or portion of an organism's neurally-stored information, identified using the abstraction system of the observer, whose instantiation depends critically on causation by prior instantiation of the same memory item in one or more organism's nervous systems.' By contrast, heteroderivative mnemons are somehow produced in response to an attempted but unsuccessful transmission of a different mnemon, or as I think Lynch might say, whose instantiation is not a

consequence of any such causation. Lynch is therefore postulating that only a subset of Dawkins B memes should be considered as true memes. But this is not a fixed subset, as mnemons can be *simultaneously* memes *and* not memes. For instance, in Lynch's system, whenever I utter the phrase `Napoleon died in 1821' (N1821) to my student, the following events occur.

Two mnemons are involved. The first is the mnemon N1821, as a basic unit of information, and the second is `belief in N1821'. Our mnemon, N1821, which I believe and therefore, Lynch would say, constitutes a mnemon `belief in N1821', is a meme *only* if my student believes it, since only then does a homoderivative replication event take place. If I have a lecture theatre full of students of varying degrees of scepticism and gullibility, and I utter N1821 (which *I* believe), it is simultaneously a meme *and* not a meme depending on the belief reactions of the students. However, its alter ego, the simple mnemon `awareness of the concept N1821' *is* always a meme, provided the students comprehend it, regardless of whether or not any of them believe it.

Aside from the awkwardness of N1821 being simultaneously a meme and not a meme, the problem for any empirical memeticist is that we cannot observe if belief has been replicated or not, so we cannot in any case decide in Lynch's system if the mnemon `belief in N1821' is a meme or not. Besides, Lynch's system still presents the same problems of defining when and how a mnemon is instantiated in its host. Since mnemons can only be inferred on the basis of behaviour or artefacts, adoption of these complicated neologisms does not clarify any of the issues raised here (and surely the last thing memetics needs is yet more neologisms). Lynch also presents `a symbolic calculus of mnemon conjugations and replication events', but this is not usable in the elucidation of any of the problems here presented, since again the invalid assumptions are made at the outset that individuals *have* mnemons, and that belief states are transmissible mnemons.

8.6 That I Contradict Statements Made In My Earlier Work

This is true, so I ought to be specific about how much backtracking I wish to do.

In Gatherer (1998), p.209-210), I state:

> *`Rather than mind viruses, religious memes, like scientific ones, may be considered as mind symbionts, replicating themselves through a positive contribution to the well-being of (most of) those who carry them.'*

I am now considerably less comfortable with the idea of individuals *carrying* any memes. Religious beliefs may be beneficial to (most of) those who profess them, but this is a psychological phenomenon rather than a memetic one. The whole mind-virus/symbiont controversy is based on an incorrect meme-host duality, as described in the present article. I still stand by my view of religions as `large, integrated complexes of memes' (Gatherer 1998, p.205), but I should now say that beliefs are not among those memes. The memes in question are the observable, copiable aspects of religion, such as prayers, rituals, artefacts, liturgies, dogmas etc.

In an earlier article (Gatherer 1997b), I discuss the merits of memetic diversity, and the problems arising from its global depletion. At the time I intended beliefs to be considered as part of such memetic diversity. However, this argument loses no force even when belief is left out of the issue. In fact, that much was already implicit in the original. We may preserve the memes of indigenous cultures, their practices, artefacts and general way of life, but we cannot preserve their beliefs. I can dress like an Amazonian shaman, dance like him and perform his rituals, but I *cannot*, try as I might, *believe* the same things as he does; belief is simply not transmissible.

By the time I came to write Gatherer (1997c), I was already effectively considering memes as observable artefacts (music, in the context of that paper). I had not then consciously come to reject a mentalistic notion of memes, but perhaps the rigours of considering a real cultural system were beginning to push me in that direction.

---

9 And Finally....Yet Another Meme Definition

One who spends so much effort criticising the meme definitions of others, ought to present a clear target for his own critics. The following is therefore offered:

> ***Meme**: an observable cultural phenomenon, such as a behaviour, artefact or an objective piece of information, which is copied, imitated or learned, and thus may replicate within a cultural system. Objective information includes instructions, norms, rules, institutions and social practices provided they are observable.*

That will suffice for a bald definition. By way of supplementary clarification, one might say that this definition excludes behaviours which are genetically determined (which are left to the sociobiologists), although a behaviour which is genetically determined in one individual may be copied by another individual who is not so determined, and thus become a culturally propagated meme (eg. incest avoidance, see Durham 1991, Chapter 6). It also excludes those aspects of culture which are not observable, and not readily transmissible, such as beliefs (which are left to the cognitive psychologists). It also excludes behaviours which are learned via trial-and-error or Skinnerian reinforcement although, as above, once such behaviours are observable, they can be copied and thereby become memes. Individuals may be, and usually are, intimately involved in the production, transmission and copying of memes, but these memes cannot be said to belong to such individuals, or to parasitise such individuals.

I do not claim any originality for this definition, because it is essentially identical to that previously given by Benzon (1996):

> *`.... I suggest that we regard the whole of physical culture as .... [memes]: the pots and knives, the looms and cured hides, the utterances and written words, the ploughshares and transistors, the songs and painted images, the tents and stone fortifications, the dances and sculpted figures, all of it. For these are the things which people exchange with one another, through which they interact with one another. They can be counted and classified and variously studied.' (Benzon 1996, p.323)*

One might ask: what difference would such a memetics have to already existing fields such as diffusion studies or social contagion theory? The answer would be: very little, except perhaps that a Darwinian streak would be more prominent; it would be an empirically-based evolutionary contagion/diffusion theory, which might *eventually* tell us something about the evolution of culture as a whole.

But this would be no bad thing. Memetics as it currently stands is a strange beast, an unrealistic psycho-epidemiology, laden with neologisms, neither theory of mind nor theory of culture. To make it into a science, we need to move forward into territory already occupied by diffusion sociologists and social psychologists. *We are not pioneers* (despite the 'new science of memetics' rhetoric that occasionally surfaces), but just *a new wave of settlers*. To survive in our new environment we need to learn from those who have been living there at least since the days of Gabriel Tarde at the turn of the century. The first step must be to shed the unnecessary baggage of `thought contagion', `memory abstractions' and meme-host duality. This is not just necessary in order to retain the goodwill and respect of our new neighbours,

but for the sake of intellectual survival.

Acknowledgements

I thank the members of the Bristol Meme Lab for kindly inviting me to visit them, and for many helpful comments. Similar thanks go to Stephen Bentley and his evolutionary discussion group at University College London. My colleagues Nick Owen (Liverpool Institute of the Performing Arts) and Holly Smith (School of Human Sciences, Liverpool JMU) have also suffered having my ideas bounced off them. None of these people, of course, necessarily agree with any of the above.

References

Ball, J.A. (1985) Memes as replicators. *Ethology and Sociobiology* **5**, 145-161.

Barbrook, R. (1996) Never mind the cyberbollocks... A contribution to *Memesis* http://web.aec.at/meme/symp/panel/msg00076.html

Benzon, W. (1996) Culture as an evolutionary arena. *Journal of Social and Evolutionary Systems* **19**, 321-362. http://www.newsavanna.com/wlb/CE/Arena/Arena00.shtml

Brown, L. (undated, probably 1964) *A Bibliography on Spatial Diffusion - with Special Emphasis on Methodology and Theory.* Department of Geography, Northwestern University, Evanston, Illinois.

Cavalli-Sforza, L.L. and Feldman. M.W. (1981) *Cultural Transmission and Evolution: A Quantitative Approach*. Princeton University Press: Princeton.

Cavalli-Sforza, L.L., Feldman, M.W., Chen, K.H. and Dornbusch, S.M. (1982) Theory and observation in cultural transmission. *Science* **218**, 19-27.

Cavalli-Sforza, L.L. and Feldman, M.W. (1983) Cultural versus genetic adaptation. *Proc. Natl. Acad. Sci. USA*, **80**, 4993-4996.

Cavalli-Sforza, L.L. (1986) Cultural evolution. *American Zoologist* **26**, 845-855.

Cohen, D. (1979) *J.B.Watson: the Founder of Behaviourism: a Biography*. Routledge, London.

Crandall, C.S. (1988) Social contagion of binge eating. *Journal of Personality and Social Psychology* **55**, 588-598.

Dawkins, R. (1976) *The Selfish Gene*. Paperback edition 1978. Granada Paladin: London.

Dawkins, R. (1982) *The Extended Phenotype*. Paperback edition 1983. Oxford University Press: Oxford.

Dawkins, R. (1993) Viruses of the mind. In: *Dennett and his Critics*. Dahlbom, B. (Ed.) Blackwell: Oxford. pp.13-27.

Dennett, D.C. (1990) Memes and the exploitation of the imagination. *Journal of Aesthetics and Art*

Criticism **48**, 127-135. http://www.tufts.edu/as/cogstud/papers/memeimag.htm

Dennett, D.C. (1991) *Consciousness Explained*. Little Brown: Boston.

Doran, J. (1998) Simulating collective misbelief. *Journal of Artificial Societies and Social Simulation* **1**, no. 1. http://www.soc.surrey.ac.uk/JASSS/1/1/3.html

Durham, W.H. (1991) *Coevolution: Genes, Culture and Human Diversity*. Stanford University Press: Stanford.

Erasmus, C.J. (1952) Changing folk beliefs and the relativity of empirical knowledge. *Southwestern Journal of Anthropology* **8**, 411-428.

Feldman, M.W. and Cavalli-Sforza, L.L. (1984) Cultural and biological evolutionary processes - gene culture disequilibrium. *Proc. Natl. Acad. Sci. USA*, **81**, 1604-1607.

Feldman, M.W., Cavalli-Sforza, L.L. and Peck, J.R. (1985) Gene culture coevolution - models for the evolution of altruism with cultural transmission. *Proc. Natl. Acad. Sci. USA*, **82**, 5814-5818.

Ficken, M.S. and Popp, J.W. (1995) Long-term persistence of a culturally transmitted vocalisation of the black-capped chickadee. *Animal Behaviour* **50**, 683-693.

FlyBase (1998) FlyBase Gene Report. FlyBase: a Database of the Drosophila Genome. Department of Biology, Indiana University, Jordan Hall, Bloomington, IN 47405. http://flybase.bio.indiana.edu:82/exec+:-vtext/html_-dR58262235-59314610-/.data/genes/genes.txt:/

Freedman, J.L. and Perlick, D. (1979) Crowding, contagion, and laughter. *Journal of Experimental Social Psychology* **15**, 295-303.

Freud, S. (1959) *Group Psychology and the Analysis of the Ego*. Trans. and ed. J. Strachey. Norton: London.

Gabora, L. M. (1997) The origin and evolution of culture and creativity *Journal of Memetics - Evolutionary Models of Information Transmission* **1**. http://cfpm.org/jom-emit/1997/vol1/gabora_l.html

Gatherer, D. (1997a) Macromemetics: towards a framework for the re-unification of philosophy. *Journal of Memetics - Evolutionary Models of Information Transmission* 1. http://cfpm.org/jom-emit/1997/vol1/gatherer_dg.html

Gatherer, D. (1997b) Feyerabend, Dawkins and the politics of cultural diversity. *Anarchist Studies* **5**, 23-43.

Gatherer, D. (1997c) The evolution of music: a comparison of Darwinian and dialectical methods. *Journal of Social and Evolutionary Systems* **20**, 75-92.

Gatherer, D. (1998) Meme pools, World 3 and Averroës's vision of immortality. *Zygon: Journal of Religion and Science* **33**, 203-219.

Gould, S.J. (1996) *Life's Grandeur: the Spread of Excellence from Plato to Darwin*. Cape, London.

Guglielmino, C.R., Viganotti, C., Hewlett, B. & Cavalli-Sforza, L.L. (1995) Cultural variation in Africa - role of mechanisms of transmission and adaptation. *Proc. Natl. Acad. Sci. USA*, **92**, 7585-7589.

Hamilton, J.D. and Hamilton, L.C. (1981) Models of social contagion. *Journal of Mathematical Sociology* **8**, 133-160.

Harms, W. (1996) Cultural evolution and the variable phenotype. *Biology and Philosophy* **11**, 357-375.

Hawley, F. (1946) The role of Pueblo social organization and the dissemination of Catholicism. *Applied Anthropology* **48**, 407-415.

Hewlett, B.S. and Cavalli-Sforza, L.L. (1986) Cultural transmission among Aka pygmies. *American Anthropologist* **88**, 922-934.

Laland, K.N. (1992) A theoretical investigation of the role of social transmission in evolution. *Ethology and Sociobiology* **13**, 87-113.

Lefkowitz, M., Blake, R.R. and Mouton, J.S. (1955) Status factors in pedestrian violation of traffic signals. *Journal of Abnormal and Social Psychology* **51**, 704-706.

Levy, D.A. and Nail, P.R. (1993) Contagion: A theoretical and empirical review and reconceptualization. *Genetic, Social and General Psychology Monographs* **119**, 235-285.

Lumsden, C.J. and Wilson, E.O. (1985) The relation between biological and cultural evolution. *Journal of Social and Biological Structures*. **8**, 343-359.

Lynch, A. (1996a) *Thought Contagion: How Belief Spreads through Society: The New Science of Memes*. Basic Books, New York.

Lynch, A. (1996b) Contribution to `Memes: Self-replicants or Mysticism.' *Wired* debate October 1996. http://www.wired.com/braintennis/96/43/bt0a.html

Lynch, A. (1998) Units, events and dynamics in memetic evolution. *Journal of Memetics - Evolutionary Models of Information Transmission* **2** http://cfpm.org/jom-emit/1998/vol2/lynch_a.html

Moss, P.D. and McEvedy, C.P. (1966) An epidemic of overbreathing among schoolgirls. *British Medical Journal* **1966/2**, 1295-1300.

Pennebaker, J.W. (1980) Perceptual and environmental determinants of coughing. *Basic and Applied Social Psychology* **1**, 83-91.

Phillips, D.P. (1974) The influence of suggestion on suicide: substantive and theoretical implications of the Werther Effect. *American Sociological Review* **39**, 340-354.

Pinker, S. (1998) *How the Mind Works*. Allen Lane The Penguin Press, London.

Preti, A. and Miotto, P. (1997) Creativity, evolution and mental illnesses. *Journal of Memetics - Evolutionary Models of Information Transmission*, **1**. http://cfpm.org/jom-emit/1997/vol1/preti_a&miotto_p.html

Rapaport, A. and Rebhun, L.I. (1952) On the mathematical theory of rumor spread. *Bulletin of Mathematical Biophysics* **14**, 375-384.

Rashevsky, N. (1951) Mathematical biosociology of beliefs and prejudices. *Bulletin of Mathematical Biophysics* **13**, 289-302.

Rashevsky, N. (1952) Prolegomena to a dynamics of ideologies. *Bulletin of Mathematical Biophysics* **14**, 95-118.

Richerson, P.J. and Boyd, R. (1978) A dual inheritance model of the human evolutionary process I: basic postulates and a simple model. *Journal of Social and Biological Structures* **1**, 127-154.

Rogers, E.M. and Shoemaker, F.F. (1971) *Communication of Innovations. A Cross-cultural Approach*. 2nd ed. Free Press: New York.

Rose, S., Kamin, L.J. and Lewontin, R.C (1984) *Not in Our Genes: Biology, Ideology and Human Nature*. Penguin, London.

Sigelman, C.K. and Sigelman, L. (1976) Authority and conformity: violation of a traffic regulation. *Journal of Social Psychology* **100**, 35-43.

Skinner, B.F. (1972) *Beyond Freedom and Dignity*. Cape: London.

Speel, H-C. (1997) A memetic analysis of policy making. *Journal of Memetics - Evolutionary Models of Information Transmission*, **1**. http://cfpm.org/jom-emit/1997/vol1/speel_h-c.html

Stahl, S.M. and Lebedun, M. (1974) Mystery gas: an analysis of mass hysteria. *Journal of Health and Social Behaviour* **15**, 44-50.

Takahasi, K. (1998) Evolution of transmission bias in cultural inheritance. *Journal of Theoretical Biology* **190**, 147-159.

Wheeler, L. and Caggiula, A.R. (1966) The contagion of aggression. *Journal of Experimental Social Psychology* **2**, 1-10.

Wilkins, J. S. (1998) What's in a meme? Reflections from the perspective of the history and philosophy of evolutionary biology. *Journal of Memetics - Evolutionary Models of Information Transmission*, **2**. http://cfpm.org/jom-emit/1998/vol2/wilkins_js.html

de Winter, K.W. (1984) Biological and cultural evolution: different manifestations of the same principle. A systems-theoretical approach. *Journal of Human Evolution* **13**, 61-70.

Back to Issue 2 Volume 2