Working Draft – Comments Welcome!

Abstract Cultural Replicators

Justin C. Fisher
University of Arizona

In recent years meme-theoretic approaches to understanding cultural change have received an increasing amount of popular attention. These approaches propose that we might understand cultural change as being driven, at least in large part, by a competition between socially transmitted replicating entities called 'memes'. One outspoken critic of these approaches is the anthropologist Dan Sperber. In this paper, I defend a general meme-theoretic approach against several criticisms that Sperber raises. This defense helps bring to light some ways in which a defensible meme theory could (and perhaps should) be developed.

Motivating Meme Theory

I will open by offering just a few words regarding why it is that I find meme-theoretic approaches attractive. Since Richard Dawkins has put this quite well before me (though perhaps without recognizing the full impact of what he was saying), I will begin by borrowing some words of his.

> Darwin's 'survival of the fittest' is really a special case of a more general law of *survival of the stable*. The universe is populated by stable things. A stable thing is a collection of atoms that is permanent enough or common enough to deserve a name. It may be a unique collection of atoms, such as the Matterhorn, [...] or it may be a *class* of entities, such as raindrops, that come into existence at a sufficiently high rate [...] even if any one of them is short-lived. The things we see around us that we think of as needing explanation—rocks, galaxies, ocean waves—are all, to a greater or lesser extent, stable patterns of atoms. (Dawkins 1976, pg 12).

Science is in the business of finding recurring patterns in the world. Wherever there is an explainable recurring pattern, there is, at least at some level of abstraction, an entity potentially worth naming—specifically, whatever it is that instantiates that pattern.

Dawkins' *The Selfish Gene* describes one special way in which certain complex patterns might sustain their recurrence over long periods of time: each instantiation might be causally disposed, under certain reasonably common circumstances, to bring it about that another instantiation of the same type comes into existence. Dawkins argues, rightly, that very interesting selectional consequences might naturally arise from a system of entities—he calls them *replicators*—that depend upon this sort of process for their continued existence. The classic example is, of course, the *gene*. However, Dawkins argues, the general principles underlying genetic selection apply quite generally to all replicators. To illustrate the possibility of other replicators, he coined the word *meme* as a name for any *cultural* entity whose repeated recurrence depends upon a similar process—i.e., memes are (generally quite abstract) entities that are causally disposed, under certain reasonably common circumstances, to bring into existence new entities that are relevantly like themselves. What distinguishes *memes* from *genes* is their manner of replication—memes reproduce themselves through people's social interactions, while genes reproduce themselves through DNA synthesis.

As I understand memes, it is clear that very many memes exist. It is clear that many recurring (though often quite abstract) patterns in human behavior are socially transmitted. 'Meme' is just a general term for any such pattern.

However, just because memes exist, it does not necessarily follow that they will be all that interesting to study. It might be, for example, that memes are too short-lived or mutable to sustain any noteworthy higher-order processes. On the other hand, though, it is clear that many interesting cultural patterns do occur, and are worth studying—such patterns underlie most work in linguistics, anthropology, social psychology, economics and most of the social sciences and humanities. I suspect that wherever such interesting patterns occur, at some appropriate level of abstraction, there are memes; and that, at this level of abstraction, these patterns may be well understood in terms of a process of memetic change.

This is intended more as a motivational picture, rather than as a compelling argument. Compelling arguments for this point can come only through hard-earned empirical results. However, this motivational picture might serve—and indeed does serve for me—as a cause for optimism.

Sperber's Objections.

Let us move on to consider Dan Sperber's objections to meme theoretic approaches. Sperber thinks that it is appropriate to understand cultural change in terms of the ways that various 'mental representations' bring other 'mental representations' into existence, and the ways in which the distribution of these representations changes over time. By itself, this is entirely consistent with many meme-theoretic approaches. However, Sperber argues that this process of change cannot generally be well-understood in terms of a replication process, as meme theory would suggest. As he puts it, “My two basic points [...] have been (1) that representations don’t in general replicate in the process of transmission, they transform; and (2) that they transform as the result of a constructive cognitive process. Replication when it truly occurs, is best seen as a limiting case of zero transformation.” (Sperber 1996, pg 101).

Sperber (1996, 2000) offers several interesting examples that he thinks illustrate that cultural change is best understood as being a transformational process—rather than as being a replicative process, as meme theory would suggest. Let us begin by considering one of these examples carefully.

> Imagine a population of items that are individually capable of begetting descendants, and that have a limited life-span. Let us imagine that these items come in 100 types, with relationships of similarity among the types such that we may represent the space of possibilities by means of a 10 by 10 matrix [...]. Imagine some initial stage (which might, e.g., have been experimentally contrived) at which we have a random distribution of, say, 10,000 items among the 100 types. Suppose we examine our population after a number of generations, and observe a different distribution [as in Fig. 1]. (Sperber, 1996, pg 109-10).

Sperber grants that one possible explanation for such a change in distribution is the obvious replicator-theoretic one—it might be that most entities beget near-identical replicas, and that the types of replicators that ended up most well-represented in the final distribution were the types of replicators that were causally disposed to beget the most (or best) replicas.

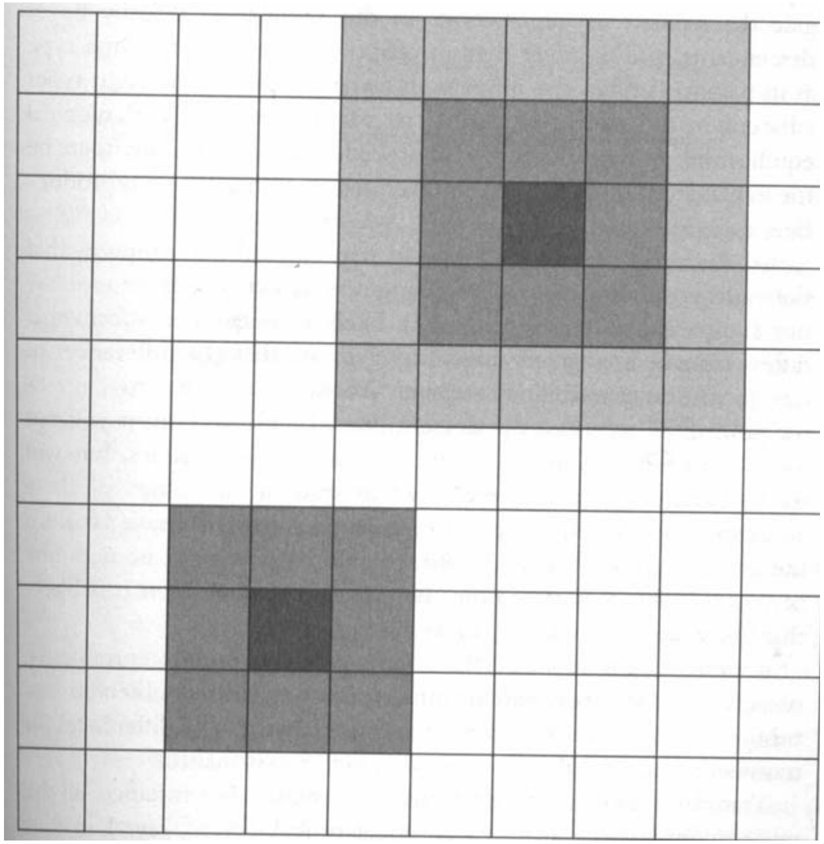

Figure 1. “The space of possibilities. After a few generations, density is greater in the two shaded areas.” (scanned from Sperber 1996, pg 109.)

However, Sperber suggests that another explanation might equally account for these results, and in a way that provides a more plausible model for cultural change. He illustrates this alternate manner of explanation with an example that seems as though it clearly excludes replicator-theoretic explanation. "Suppose, however, that we [...] discover that an offspring is *never* of the same type as its parent! Rather, the offspring is always of one of the eight types adjacent in the matrix to that of its parent [see Figure 2]." (Sperber 1996, pg. 110).

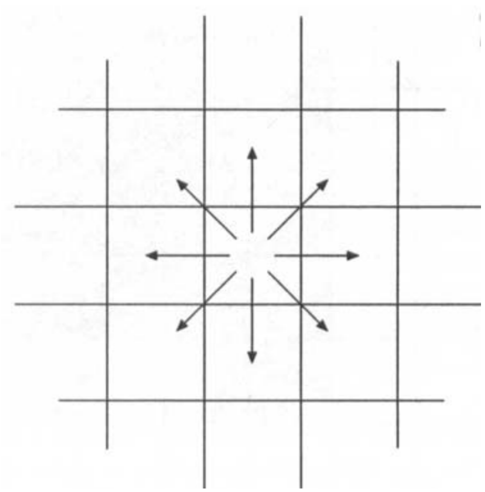

Figure 2. "An item begets descendants of neighboring types" (scanned from Sperber 1996, pg 111).

Sperber suggests that the resulting distribution might then be explained by biases in the begetting process; biases that make certain types of offspring disproportionately likely to arise in each begetting as compared to the other possible alternatives. Sperber suggests that we might understand what is going on here in terms of 'attractors'. As I will use the term, an 'attractor' is any maximal closed[1] set of similar types (i.e., ones that are contiguous in the matrix) that share such a disproportionate advantage over types that 'neighbor' the set. Alternatively, one might say that 'the attractor' is the single type at the 'center' of such a set, or the single type in the set that enjoys the greatest advantage. It's not clear to me which, if any, of these precise notions Sperber has in mind.[2] What is important, though, is that a transformation process might be

[1] By *maximal* I mean that adding any 'neighboring' types to the set would reduce the overall disproportionate advantage that distinguishes the set. By *closed* I mean that the set does not completely 'surround' any (small) set of types that are not in the given set. This definition is meant only as a heuristic—as a way for us to get an abstract handle of what is going on in a complex system. It is fine if there are cases in which it is unclear (or to a certain degree arbitrary) what we might call 'attractors'.

[2] Sperber seems to have in mind one of these latter 'single-point' formulations. However, I dislike them because they seem to rule out extended attractors. E.g., a system might have three very strong attractors that together display the appearance of the percent-sign '%'. Any entity near one of these three extended attractors would be likely to have a descent-line that stays stably within that attractor for some time, even as the descent line moves quite far from any particular point in that attractor.

biased in such a way that it will have two distinct attractors, and that the characteristics of such a process might explain the distributions that result. I have run computer simulations (described below) that verify the possibility of results much as Sperber describes. However, I will show that such computer simulations may also reveal results which weigh against Sperber's ultimate conclusion.

Sperber suggests that this example poses serious concerns about the viability of a meme-theoretic approach. On the surface, this example seemed to be as good a case as we might hope for replicator-theoretic analysis. It would seem that the relevant entities are clearly identifiable, and so too are their relations of descent. However, Sperber suggests, a replicator-theoretic analysis *can't* explain what's going on in this case. "Ecological equilibrium among differently endowed replicators cannot then be the explanation since this is a transformative, rather than a reproductive, descent system" (Sperber 1996, pg 110).

By itself, this doesn't weigh against meme-theory—no meme-theorist is committed to saying that every single dynamical system is well understandable in terms of replicators.[3] Sperber's objection gets its teeth, however, from his plausible observation that most cases of cultural transmission probably do involve at least a certain degree of transformation. Entities acquired through cultural transmission are often somewhat like the originals, but quite often they also differ from the originals in noteworthy ways. Hence, the meme-theorist can't be cavalier about models like this—for it is plausible that these models do reflect certain important characteristics that are present also in the case of cultural change. If Sperber is right that a replicator-theoretic approach is ill-suited to models like this, then this does bode somewhat poorly for the prospects of meme theory.

However, I don't think Sperber is right to claim this. To see why this is, though, we will need to explore carefully what is meant by the term 'replicator.' As we will see, Sperber has some rather specific ideas about what a replicator is. However, not all of these ideas are ones we should accept.

What is a Replicator?

Dawkins introduced the term 'replicator' as a way of referring to any entity that has "the extraordinary property of

[3] Though, as the first section of this paper may have indicated, I actually would be inclined to defend a somewhat qualified version of this claim.

being able to make copies of itself" (Dawkins 1976, pg 15). He suggests that more successful replicators will tend to measure more highly on three important dimensions: *longevity* (the capacity to survive intact up to the point of replication), *fecundity* (the disposition to create numerous viable copies of oneself) and *copying-fidelity* (the tendency to create copies that are similar to oneself in relevant ways).

However, this is probably not a precise enough specification to arbitrate the present issues. What's more, as Sperber suggests, there is a significant danger in blindly counting as a replicator *any* entity that causes the production of things like itself. For example, consider contagious laughter or the ways that dogs' barking in the night inspires other dogs to bark. In each of these we may trace a causal chain of tokens causing the existence of similar tokens. However, it is quite clear that the successive tokens are not full-fledged replications of the previous tokens.[4] To capture this intuitive idea, Sperber suggests that the following constitute at least necessary conditions on replication:

> For B to be a replication of A,
> (1) B must be caused by A (together with background conditions),
> (2) B must be similar in relevant respects to A, and
> (3) The process that generates B must obtain the information that makes B similar to A from A. (Sperber 2000, pg 169).

How should the meme theorist respond to these conditions? Condition (1) is uncontroversial. So long as we are careful with the tricky word 'relevant', Condition (2) should also pose no problems. Condition (3) is much harder to evaluate, largely because it is not entirely clear what Sperber means by 'obtaining information', nor how 'obtaining information' might serve to 'make' one thing similar to another. Let's consider what Sperber offers as a "[a]nother way to express this third condition":

> (3') B must inherit from A the properties that make it relevantly similar to A. (ibid.)

Now, (3') sounds plausible, but it seems much too vague to be of much use. It's also unclear how (3') could be equivalent to (3), because it's not at all clear that (3') says anything at all about 'information'. However, Sperber clearly does want to

[4] Regarding a similar case, below, I argue that there is something that is replicated along these chains. In the present cases, that something is the determinate (binary) value of whether or not each successive individual is laughing/barking. Of course, the limited form of replication that occurs in these cases is unlikely to be of much interest.

make much of the informational aspect of (3) in his evaluation of some example cases (as we will see in a moment). Hence, I think we must stick to considering (3) as we evaluate his argument.

One reasonable approach to understanding information derives from the classic work of Shannon and Wiener[5]. This approach has it that information is anything that might serve to eliminate certain states of affairs that might otherwise have been possible. For example, information about the world serves to reduce the number of possible ways the world might be, for all some subject knows. With respect to a replication (or pseudo-replication) process, we may view informational inputs as being things that constrain which sorts of outputs are possible from the process. On a very strong reading of (3), the only information that can constrain the final shape of a replication is information about the original. This strong reading of (3) denies that any antecedent cognitive information might play any vital role in a replication process. I think this strong reading does capture the ways that Sperber applies (3) to certain example cases, as we will see in a moment. However, I think this strong reading is not plausibly true.

For, this strong reading of (3) eliminates far too many replication processes. Take for example, DNA synthesis. DNA synthesis typically involves the joining together of numerous pre-constructed nucleotides. Which pre-constructed nucleotides are available constitutes one constraint upon the possible outcomes of this process. On a very broad reading of 'information,' then, these available nucleotides constitute some information regarding what the final state of the process might be. The case becomes even worse when we consider enzymes whose function it is to repair copying imperfections in replicated DNA. It is clear that these enzymes constrain the possible outcomes of the process—hence that they constitute information. Yet clearly, DNA synthesis is still a paradigm case of replication, despite the intrusion of these other sorts of information. The strong reading of (3) is too strong.

How might we weaken (3) to make it more plausible? I think the intuitive idea motivating (3) is not that other forms of information *can't* be involved in a replication process—it is just that any replication process must be appropriately sensitive to the properties of the thing that is being replicated. I propose that this is well captured in the following formulation:

[5] For a good overview, see Dretske 1980.

> (3'') If B is a replication of A, then A's having the relevant properties must play an appropriate *causal role* in B's coming to have relevantly similar properties.

This condition serves to rule out cases in which B comes to have the relevant properties in some way that is quite independent of A's having those properties (or similar ones). It is worth noting that my condition (3'') is closely related to a condition that Ruth Millikan (1984) proposes in defining her own version of the replication:

> (3''') [H]*ad A been different* with respect to [the relevant properties], as a result, *B would have differed accordingly*. (Millikan 1984, pg 20).

Indeed, on certain counter-factual understandings of causation, my (3'') would entail Millikan's (3''').

Millikan's work also suggests one useful way of thinking of the 'relevant' properties that must be similar between originals and their replications. She suggests that we should think of each such property as being one among a number of mutually exclusive *determinate* values that some more general *determinable* property might take (ibid.). For example, *color* is a *determinable* property; *red* and *green* are *determinate* properties. As another example, it is *determinable* which base-pair occupies a particular chromosomal locus—the *determinate* alternatives are the four base-pairs. We may read (3''') as requiring that a replication process be likely to convey the *determinate* values of certain *determinable* properties from original to replication.

How well do these new formulations capture what Sperber had in mind? It seems, at least, that these formulations do well to account for the example that Sperber offers to motivate (3). "Ten sound-recorders with the same repertoire of melodies in each have been fixed so that they are activated by the sound of the last five bars of any melody in their repertoire, and then play this very melody" (Sperber 2000, pg. 169). The recorders are arranged in a queue, so that each one triggers the next to play the same melody that it played. On first appearances, it would be dangerously easy to judge that each recorder was replicating the full melody—note-by-note—that the previous recorder played. I think Sperber is right that this initial judgment would be mistaken. The reason for this is that the process that brings about the numerous similarities between successive melodies does not causally depend upon the preceding melodies' displaying (very many of) those properties.

It is worth pointing out, though, that *something* is replicated in this chain—namely the fact that it is Melody #23 that each recorder plays, rather than some other melody. (I.e., that *playing Melody #23* is the determinate value of the determinable *which-melody-is-being-played*.) If we take similarity with respect to *which-melody-is-being-played* to constitute the relevant similarity, then we have something that should meet the conditions for replication, afterall. The presence of the property *playing Melody #23* in each subsequent recorder causally depends in the right ways on the presence of that property (or one quite similar to it) in the preceding recorder. Hence, while Sperber is right to insist that, in the case he describes, the melodies are not being fully replicated (note-by-note), he would be wrong to hold that this case doesn't involve replication at all.[6]

[6] One potential confusion in this case involves the possible equivocation between *collective* and *distributive* readings of 'melody'. E.g., it would be a mistake to equate 'The melody was beautiful' with 'Each note in the melody was beautiful.' It would be a similar mistake to equate 'The melody was replicated' with 'Each note in the melody was replicated.' The former (on certain collective readings) might be true, even while the latter (which forces a distributive reading) is false.

Applying the Three Conditions

While these suggested reformulations of (3) do seem to deliver the result Sperber wanted regarding his chain of recorders, I think they depart from his thinking regarding other cases. But whenever (3) and (3'') disagree about particular cases, I think that my (3'') is more plausibly correct. We have already seen that Sperber's (3) has trouble handling actual cases of DNA replication. In a moment we will see how Sperber applies (3) to another case, one proposed by Dawkins (1999). But first, let's consider a similar case that will help to make clear many of the issues that also arise in the Dawkins/Sperber case.

Consider the following two photo-copy machines. Machine #1 is an ordinary photo-copy machine. It is intuitively clear that such a machine is capable of replicating printed pages. Machine #2 begins by running the original sheet through a somewhat low-resolution scanning device. The scanned image is fed into an Optical Character Recognition program. This program employs rich a rich database encoding information about grammar, word-spelling, and word frequencies. With the help of this database, this program infers, quite reliably at least for pages printed in English, what the original page must have looked like, and it directs a laser printer to print off a page that looks like that. It is intuitively clear to me that what Machine #2 produces is also a replication of the original page. At least it replicates the values of very many of the determinable properties of the original page. (Of course there are uncountably many determinable properties like thickness-of-paper or chemical properties that neither machine aims to replicate.) This seems to me to be another counterexample to Sperber's condition (3). Machine #2 is quite capable of making replications, despite the vital role that antecedent information plays in this process. This weighs in favor of my alternate formulation (3'').

With this in mind, let us consider the case proposed by Richard Dawkins (1999). Dawkins suggests that we imagine subjects being shown a sketch of a Chinese junk for a few minutes. Then, each subject attempts, by memory, to produce a replication of this sketch. Dawkins predicts, quite plausibly, that many subjects would perform quite poorly at this task. Imagine that the replica one subject produces is then given to another person to replicate, and that replica to a further person, and so on. Dawkins predicts, plausibly, that not far down the chain, the form of the original Chinese junk will become quite unrecognizable. (Well, it may still be junk, but now in a different sense.) Now, imagine a second case. In this case, subjects are allowed to watch the process by which an origami Chinese junk is folded. Dawkins predicts that most subjects would produce significantly better replicas in this second case. In a chain of replicas made this way, imperfections (e.g., poorly-centered folding) that occur at one stage are much less likely to be passed along.

Susan Blackmore (1999) calls this second sort of process copy-the-instructions, and the first sort copy-the-product. It is intuitively clear that, at least in many cases, copy-the-instructions is likely to have higher copying-fidelity than is copy-the-product. (There is an interesting analogy between instructions/product and genotype/phenotype.) Such observations have led Blackmore and Dawkins to conclude that instruction-copying may play an important role in many meme selection processes.

Against this account, Sperber objects, in watching the original junk-making, "[y]ou recognize [...] the behaviour as an imperfect realization of an intention of a familiar and regular type rather than as the perfect realization of an intention of an unfamiliar and irregular type. The instructions that you infer are, then, informed in part by what you actually observe, and in part by what you know about human intentions, and of the type of instructions typically used in origami" (Sperber 2000, pg 171). From this, Sperber infers, "The instructions are not being 'copied' in any useful sense of the term from one participant to the next" (ibid). Insofar as most cases of cultural transmission share the structure of this case, this conclusion would bode poorly for meme-theory. However, the inference to this conclusion is valid only given Sperber's original (3), which we have already seen to be unacceptable. If we instead accept the more plausible (3''), the inference no longer succeeds.

Notice the strong parallel between the Chinese junk cases and the photo-copy machines above. Machine #2 used its rich database to recognize familiar tokens when they appeared somewhat imperfectly before it. The origami-learner is doing precisely the same thing as she uses her background knowledge to recognize an imperfect folding-of-the-corners-to-the-middle. Just as Machine #2 replicated the paper it scanned, the origami learner may replicate a Chinese junk.

So, the learner copies the junk. Does she also copy the instructions? Imagine that we slightly alter the Machine #2 case, by imagining it to involve repeated cycles, each of which begins with a word-processor document stored in a computer. In the course of each cycle, the paper is printed, scanned, and run through OCR software and a spell-checker, yielding a new word-processor document. We have seen already that each successive print-out is a replication of the preceding print-out. This is because many relevant characteristics of subsequent printouts causally depend on their antecedents' having had similar characteristics. *Mutatis mutandis*, we may argue that each subsequent word-processor document is a replication of the preceding one. Many relevant characteristics of subsequent documents causally depend upon their antecedents' having had similar characteristics. Hence, in this iterated setup, Machine #2 does replicate the preceding word-processor document, which, in certain respects, is very much like copying-the-instructions.

We may make a very similar argument regarding origami-learning. Presumably, we may identify relevant (though perhaps very abstract) psychological characteristics shared by people who know how to fold a junk in the same way, and that differ from alternative characteristics possessed by people who do origami in other ways. As the learner learns to fold a junk, which relevant psychological characteristics she acquires will be causally (and counterfactually) dependent upon which of these relevant psychological characteristics the teacher has. It makes no difference whether or not other sorts of information play into this process. So long as this causal/counterfactual dependence holds, it is plausible that these psychological states are being replicated. Hence, it would seem, Sperber is wrong to insist that there is no sense in which the instructions are being 'copied' from one learner to the next.

I have argued that Sperber's (3) gives the wrong results in these cases, while my own (3'') provides more plausibly right results. Let us therefore adopt (3''), and return to Sperber's (1996) argument with which we initiated this paper.

Abstract Replicators in Sperber's Transformational Model

In this final section, we return to the argument with which we started out. Sperber asked that we imagine a population consisting of 100 types of entities. These types might be represented in a 10 by 10 matrix. Each entity is short lived, but causes the creation of other offspring entities. Sperber asks that we imagine that each offspring is of a similar type to its parent, but never of the same type. There are biases making some types of offspring more likely to be produced than are other types of offspring that a particular parent might produce. Sperber asks that we imagine that these biases are such that, after a number of generations, the population of entities becomes distributed in such a way that most living entities fall within (or near) either of two attractors.

Sperber takes this case to be one that is like the process of cultural change in at least one important way—entities cause the production of new entities that are similar but often not identical to the originals. Sperber suggests that a replicator-based approach cannot explain what is going on in this model, because none of the 100 types of entities ever serves to replicate itself. If this were indeed a culture-like model about which a replicator-based approach has nothing useful to say, then this would bode poorly for meme theory.

However, I will argue, there actually is quite a lot that a replicator-based account can say about this case. Now that we have a clearer notion of what a replicator might be, we are well prepared to notice replicators in this case—replicators that Sperber overlooks. Sperber infers that this system contains no replicators directly from his observation that none of the 100 types of entities ever produces an exact replica of itself. This inference is invalid for two reasons.

First, there are clear cases of replicators that would not meet this criterion. For example, DNA replicates by having a single unzipped strand of a DNA molecule facilitate the production not of a duplicate of itself—but instead of a complement or 'negative' of itself. Of course, this 'negative' strand will likely facilitate the production of a near-duplicate of the original 'positive' strand in a later replication event. In no instance of DNA replication does a strand produce a replica precisely like itself, yet clearly DNA does replicate. Hence, Sperber's inference can't be generally valid.

Second, and more import for present purposes, Sperber's argument relies on the suppressed premise that the 100 types of entities he specifies (together with the 'relevant' similarities he specifies between them), are the only potential replicators in the system. To justify Sperber's conclusion we'd need to rule out *all* the potential replicating entities in this system, not just the 100 possibilities that Sperber considers.

What other replicators might there be in this system? To answer this question, we might return to Dawkins' observation that science is in the business of finding 'stable' patterns in the world, even when those patterns are notable only at a rather high level of abstraction. What abstract stable patterns are there in this system? Sperber has been kind enough to focus our attention upon two of them, and to give them the name 'attractors'. Clearly, Sperber cannot be opposed to our identifying such abstracta in the system. "Note that an attractor, as I have characterized it, is an abstract, statistical construct, like a mutation rate or a transformation probability." (Sperber 1996, pg 111-12).

Recall that a successful replicator must be capable of causing, in virtue of its own possession of relevant properties, there to come into existence another entity that is similar to itself, at least with respect to those properties. In calling our attention to attractors, Sperber has pointed out some relevant (but abstract) properties that are copied with quite high fidelity from parents to offspring—namely being in the vicinity of an attractor. Sperber says almost as much himself. "If the departure point of a line of descent is far from the attractors, then it is likely that the arrival point will be near one of them. If the departure point is near an attractor, then it is likely that the whole line will stay in its vicinity." (Sperber 1996, pg 110). Effectively, Sperber is pointing out that *which attractor an entitiy is near* is a determinable property, and that its determinant value will be passed on with high fidelity to its offspring.

To see how this might work out, we may consider a computer simulation of the model Sperber describes. To generate this simulation, I assigned a weight to each of the 100 types of entity. When A and B are entity-types that might be produced from a particular transformational process (i.e., when both A and B neighbor an expecting parent-type), my model makes the following equality hold:

$$\frac{\Pr(A - is - produced)}{\Pr(B - is - produced)} = \frac{Weight(A)}{Weight(B)}$$

I assigned weights on the basis of Sperber's diagram (Fig. 1 above). The resulting system does stabilize upon a configuration quite a lot like the one Sperber proposes. (See Fig 3). Hence, I think this simulation captures reasonably well the intended structure of Sperber's model.

Now, let's consider some replicators we might identify in this picture. As a very rough first approximation, we might divide Sperber's 10x10 matrix into four equal quadrants. Now, we may consider an entity's quadrant to be a *determinable* property of the entity, with four possible *determinant* values (Quadrant I, Quadrant II, etc...). Employing my computer model, we may estimate the average copying-fidelity (with respect to this property) of entities in each of the four quadrants. These copying-fidelities are represented in

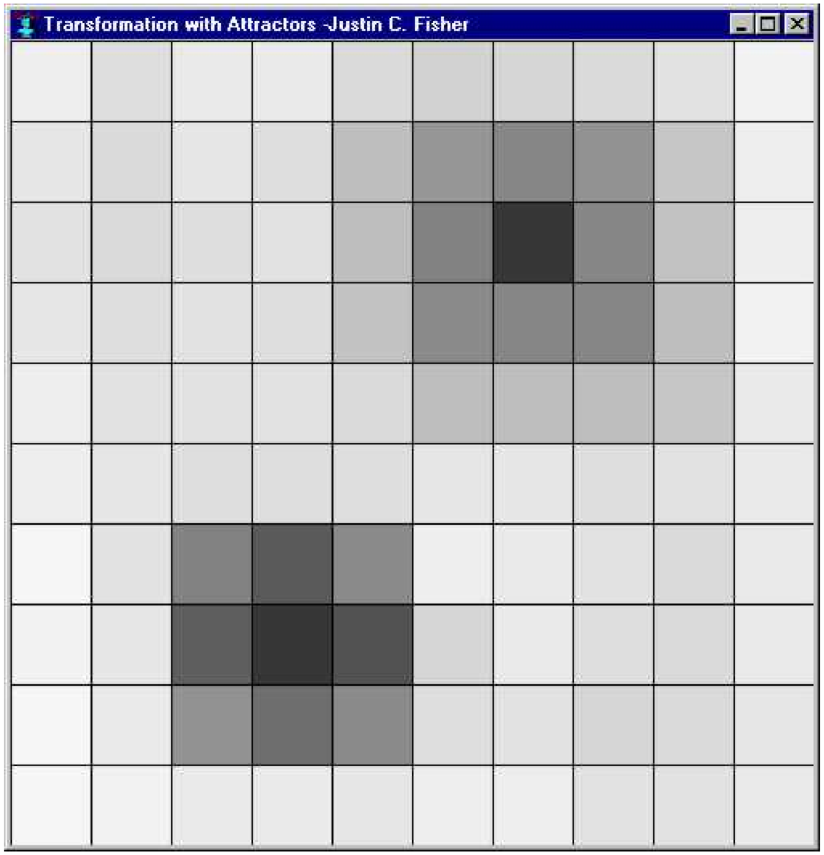

Figure 3. Population distribution after 10 generations of the descendants of 10,000 evenly distributed entities.

Figure 4, and more graphically in Figure 5. Notice, that Quadrants I and III have significantly higher copying-fidelity than do Quadrants II and IV. This reflects the presence of the 'attractors' in these quadrants. A replicator-theorist would take this fidelity difference to indicate that entities in Quadrants I and III would be more successful replicators than those in the other Quadrants, and hence that such entities would be likely to proliferate in comparison to those other entities.

Sperber would likely object that this simple four-quadrant replicator-based approach does not capture nearly all the structure that his transformational approach captures within the system. To an extent this objection is right (though I think there is much more that the replicator-theorist might say in terms of other abstract properties of this system). However, it's not clear that this objection is all that important. Compare, someone might equally object to biological explanations on the basis that they fail to capture all the details that (ideal) chemical or physical explanations of the same phenomena would capture. While correct, this objection overlooks the great value that higher levels of abstraction can have—

(II) **0.82**	(I) **0.89**
(III) **0.91**	(IV) **0.78**

Figure 4. Each quadrant's copying fidelity—the probability that a randomly chosen entity-type in that quadrant will have offspring also in that quadrant. (Estimated on the basis of 1000 trials for each entity-type.)

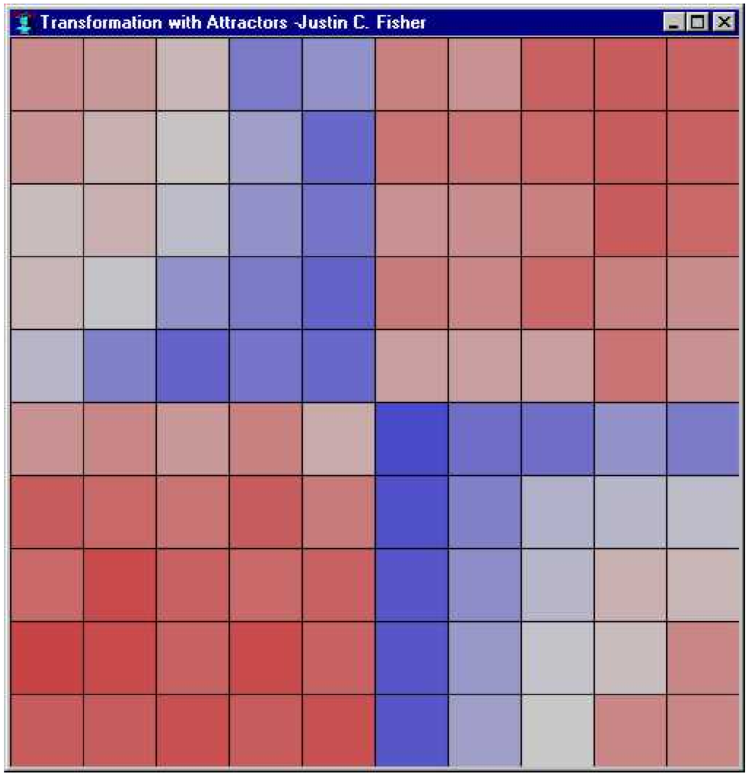

Figure 5. The probability that the 10th descendant of a given entity will be in the same quadrant as that entity. The probability is highest in the reddest squares, lowest in the bluest squares, and 0.5 in gray squares. (Estimated on the basis of 100 trials for each entity.)

especially given our practical limitations, and our desire to seek and understand recurring high-level patterns. While Sperber's transformational model might, ultimately, offer more detailed explanations of cultural phenomena—there is still much room for interesting meme-theoretic explanation too.

As a final example, let follow up on a further suggestion of Sperber's. He claims as an advantage for his transformational approach, its capacity to accounting for "[t]he multiplicity and varying number of 'parents', or sources, for the same item, which is [...] a typical aspect of cultural evolution" (Sperber 1996, pg 111).

Indeed, Sperber is right that it is not difficult to modify his model to allow for multiple parents. A simple way to do this is to imagine that each offspring entity begins life having a neutral value (i.e., as being a brain ripe for adopting a new representation). Each such offspring then observes two randomly selected parental entities. An entity-type neighboring each parent is chosen, as in the single-parent case. Then one of these two potential entity-types is chosen, again with weighted probabilities as described in the equation above. (I.e., the offspring watches two adults, provisionally acquires a transformed entity-type from each of them, and then chooses with weighted probability one of those entity-types to be its final value.)

This two-parent model produces a very interesting result, illustrated in Figure 6. After several generations, Quadrant III is densely populated, while each of the other quadrants is quite

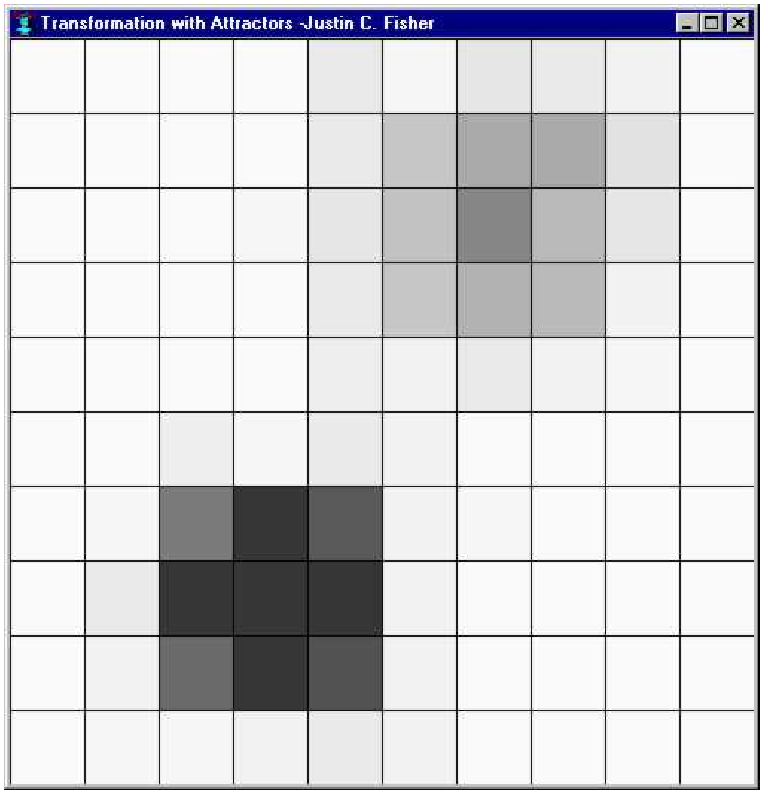

Figure 6. Population distribution after 10 generations of 2-parent reproduction, beginning with a randomly distributed population of 2000.

sparsely populated. A fine-grained transformational explanation for this result would appeal to the various transformation probabilities involved in the system—especially the fact that these probabilities are quite high (given the weightings I used) for anything in the neighborhood of the Quadrant III attractor.

We should note, however, that a replicator-theoretic approach has much to say about this case as well. Again, we may view this scenario as a competition between four types (quadrants) of replicators. We may define a replicator's fitness as the expected number of offspring produced by that replicator that will turn out to be in the same quadrant as it. This definition of fitness will be frequency dependent. The likelihood of an entity's producing/recruiting offspring that are like itself depends heavily upon the competition that that entity is likely to face from other Quadrants. When facing competition in Quadrants II and IV, entities in Quadrant I do quite well. However, these entities do much less well in the face of competition from the Quadrant III attractor. This trend is illustrated in Figure 7, which represents the 10^{th} generation of a population that began heavily skewed towards the upper half of the matrix. (It is worth noting, though, that the depicted situation is not stable—on all long runs of these simulations, a small establishment in the third quadrant gradually builds up to dominate the population, so it looks much like Figure 6.) I haven't taken the time to compute these frequency dependent fitnesses, but I am confident that Quadrant III will have the highest, followed closely by Quadrant I.

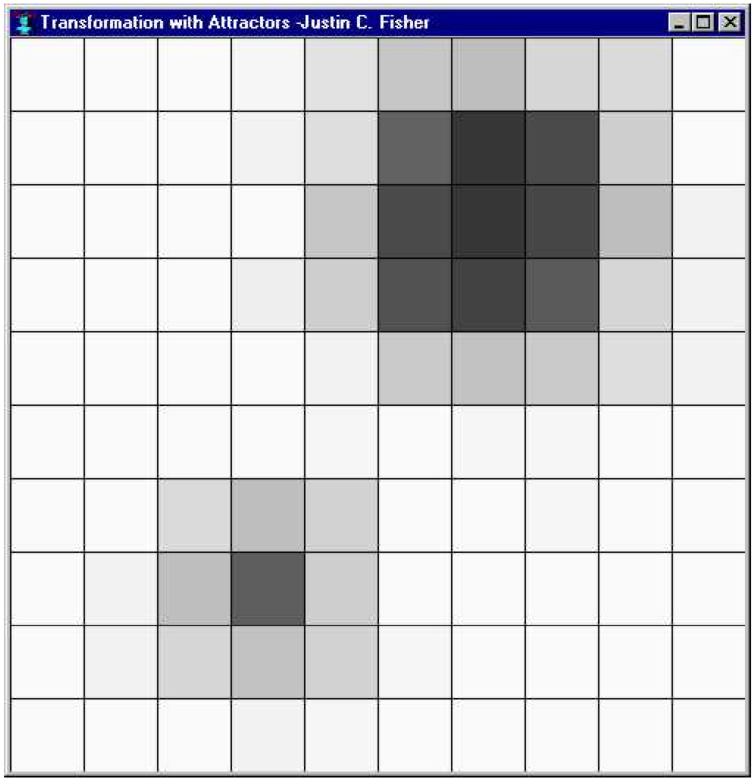

Figure 7. Population distribution after 10 generations of 2-parent reproduction, beginning with a heavily top-skewed population of 2000.

Notice, it is cases like this that are the forte of replicator-theoretic approaches. In these cases, there are clear and interesting questions that replicator-theoretic approaches can help provide answers to. Why is it that in the scenario modelled in Figure 7, Quadrant I initially experiences a population explosion? Why is it that eventually the population shifts to Quadrant III? To get helpful illuminative answers to these questions, we needn't appeal to the gritty details of Sperber's transformational model. Instead, we can give a quite good explanation at a higher level of abstraction, in terms of relative frequency-dependent fitnesses of replicators. Rather, than illustrating the need for transformational models of culture, Sperber's scenario serves to illustrate just how useful replicator-theoretic abstraction might be.

Bibliography

Blackmore, Susan. 1999. *The Meme Machine*. Oxford UP.

Dawkins, Richard. 1976. *The Selfish Gene*. New York: Oxford UP.

---, 1999. Foreword to *The Meme Machine* by Susan Blackmore. Oxford UP.

Dretske, Fred. 1980. *Knowledge and the Flow of Information*. Cambridge, MA: Bradford Books/MIT Press.

Millikan, Ruth Garrett. 1984. *Language, Thought, and Other Biological Categories*. Cambridge, MA: MIT Press.

Sperber, Dan. 1996. *Explaining Culture: A Naturalistic Approach*. Oxford: Blackwell.

--- , 2000. "An Objection to the Memetic Approach to Culture." in *Darwinizing Culture: The Status of Memetics as a Science*. Ed. Robert Aunger. Oxford UP.