

Perspectives on Imitation

17 Goals versus Memes: Explanation in the Theory of Cultural Evolution

Mark Greenberg

17.1 Introduction

The Darwinian theory of biological evolution by natural selection is an extraordinarily fruitful explanatory paradigm. When Richard Dawkins (1976/1989) introduced the idea of a “meme”—a unit of cultural transmission—his suggestion was that Darwinian explanation might also prove fertile with respect to nongenetically based cultural development. There is, of course, an obvious and commonsensical competing account of cultural change that has a strong *prima facie* plausibility in a wide range of cases: According to this *goal-based* account, it is humans’ deliberate pursuit of their conscious goals, rather than analogues of genetic mutations and natural selection, that explains why an idea or set of ideas comes into being and spreads. If a defender of a Darwinian theory of cultural evolution—*meme theory* for short—is to advance the idea that Darwinian theory is the best explanation of at least some aspects of cultural evolution, he or she must elucidate why meme theory offers explanatory advantages over competing theories, and, in particular, over a goal-based account.

What is the link to this volume’s theme? Imitation holds a central place in Darwinian theories of culture. Dawkins, for example, suggests that memes are transmitted by imitation in a broad sense and that new memes are generated by copying mistakes. The goal-based account need not deny that imitation is an important mechanism for the transmission of ideas (although it is not plausible that it is the exclusive mechanism). But the goal-based account holds that it is our goals that determine which ideas are imitated (and which changes are introduced).

I want to emphasize that my purpose is not to argue against the explanatory potential of meme theory, but to focus attention on the need to address questions of explanation. There are numerous ways in which meme

theory could offer explanatory benefits. As I will discuss, meme theorists could challenge goal-based explanations directly. For example, they could show that at least in some cases, the best explanation of the transmission of ideas is not human goals but selection in favor of "selfish" memes—ideas that are good at inducing their own replication. Or theorists could argue that even when each individual's decisions are explained by her goals, the long-term and large-scale consequences for the evolution of ideas are not what anyone wanted or intended. Thus, it might be that so-called "population-level" phenomena play a crucial role in the explanation of cultural evolution (see Boyd & Richerson, 1985, 2000). For example, the best explanation of why certain ideas are transmitted successfully from generation to generation might be that those ideas make a group more successful. A different possibility is that theorists could use meme theory to explain how human beings come to acquire their goals. Finally, once memes become an important part of the human environment, they can affect the selective pressures operating in genetic evolution; for example, genes might be selected for their effects on the transmission of memes. More generally, there is the possibility of gene-culture coevolutionary explanations.¹

In this chapter, I elaborate on the importance of addressing meme theory's explanatory power, taking Francisco Gil-White's discussion in this volume as an illustration and point of departure. Gil-White is a proponent of a Darwinian theory of cultural evolution who wants to save meme theory by distancing it from the biological theory that is its inspiration. He defends meme theory against, on the one hand, proponents who (in his view) harm it by modeling it too closely on the biological case, and, on the other, detractors who think it cannot succeed because it is too unlike the biological case. I focus on two ways in which Gil-White argues that both proponents and critics of meme theory have adhered overly closely to the biological model. He argues first, that it is wrong to require that memes be capable of high-fidelity replication, and second, that it is misguided to try to extend Dawkins's "selfish-gene" perspective to cultural evolution.

I will show that these arguments neglect the explanatory role of meme theory. First, the argument against the need for high-fidelity copying is that cumulative directional change can occur even with an extremely high mu-

1. See Cavalli-Sforza and Feldman (1973, 1981), Boyd and Richerson (1985), Blackmore (1999; vol. 2, ch. 19.12, p. 406 and vol. 1, ch. 8.3, p. 203). Feldman and Laland (1996) provide further references.

tation rate. But even if there is a process that is, broadly speaking, Darwinian in the sense that there is accumulation of change by variation and differential reproduction, it doesn't follow that Darwinian theory has anything to add to the *explanation* of that change. It is not that I am restricting the term "Darwinian" to processes that do not involve goals; rather, the point is that to the extent that humans' pursuit of their goals accounts for the course of cultural change, an appeal to Darwin is idle.

Second, the criticism of a "selfish-meme" approach overlooks the explanatory point of that approach. It is true that there are reasons other than a meme's content that can explain its proliferation; for example, the prestige of those who display it. But the selfish-meme approach does not hold merely that a meme's content is what explains its proliferation. Rather, this approach offers a potentially powerful way of challenging the view that human goals are in the driver's seat. (Indeed, the noncontent biases that Gil-White emphasizes are important in part for the very same reason.)

Before turning to a detailed discussion of these points, I want briefly to elaborate my comments about explanation. If meme theory is to pull its weight, its claim cannot be merely that ideas are transmitted with variations, that the variations accumulate gradually over time, and that the ideas that are most common now are the ones that were transmitted the most. Even if all these propositions were true, it would not follow that the Darwinian model is a helpful, or the best, explanation of cultural evolution. In particular, even if a particular cultural feature has gradually accumulated variations, it could still be the case that *the best explanation* of why the idea has developed and spread is design—planful, foresighted decisions in pursuit of people's conscious goals. For example, the best explanation of a particular technology—refrigeration, say—may be design in the service of human purposes and needs, not accumulation of copying mistakes under selective pressure. At least *prima facie*, this seems a plausible and adequate explanation of much technological development—and much else in culture as well. A special case of this account is that the best explanation of the proliferation of some ideas may be the human goal of discovering the truth. Physics and mathematics may provide examples here. The point is *not* that it is false that ideas can change by accumulation of variation under selective pressure. Rather, *even when that proposition is true*, it may not be the best explanation of what is driving the change. In sum, what a defender of meme theory needs to do is to show that an appeal to an analogue of evolution by natural selection is the best explanation of cultural evolution.

17.2 Explaining Cumulative Change

A common objection to meme theory is that the mutation, or variation, rate in the transmission of memes is too high for Darwinian evolution to take place. I will argue that it is not enough for meme theory to show that cumulative adaptive change can occur, for example, because selective pressures are extremely strong. Meme theory must also show that Darwinian theory has something to add to the explanation of the cumulative change. For example, an obvious explanation of extremely strong selection is humans' deliberate selection of ideas because of their value with respect to some independent goal. To the extent that that is the source of the "selective pressures," our goals are doing the work.

Gil-White's defense of meme theory against the objection of a high mutation rate illustrates my point.² His main argument centers on his example of Bob's tennis serve (vol. 2, ch. 16, p. 323ff). Bob has "the most attractive" serve, and everyone tries to copy it. The point is that even with a mutation rate of 1, adaptive design can accumulate. To support this conclusion, Gil-White considers a version of the example in which small errors occur every time the serve is copied, but in which the mean serve is identical to Bob's (i.e., there is no directional bias to the errors).

From the *modest* variations introduced by copying errors, a serve superior to Bob's emerges, and *this becomes the new template for us all to imitate* . . .

If we focus on the population mean, it is clear that despite the absence of replication, adaptive design accumulates under selective pressure. Moreover, the process is faster than natural selection because mutants are produced in every copying attempt. (my emphasis) (ch. 16.5.1, p. 325)

He next argues that even if there is a directional bias to the variations, design changes can still accumulate in the direction of the optimal serve. If the directional bias is toward the optimal serve, it will simply speed up evolution in that direction. If it is in the other direction, evolution toward the optimal serve can still occur as long as some variations improve on Bob's serve (again assuming everyone copies only the serve closest to the optimal serve, which displaces Bob's as the model).

I will make two points about the example. First, the example assumes what it needs to show. It is straightforward that a high mutation rate is

2. Gil-White attributes the objection primarily to Dan Sperber (1996). He does not adequately address Sperber's more important objection, based on the point that it is not true that each instance of an idea is even a low-fidelity copy of a particular ancestor idea.

consistent with cumulative directional change *if* (1) the selective pressures are strong enough and (2) the mutations are all relatively small. The tennis-serve example makes precisely these assumptions, however, and does nothing to support their plausibility. I will suggest that there is no reason to think that the corresponding assumptions will be true, or close enough to true, in a wide range of actual cases. Second, even when these assumptions are satisfied, more is needed to show that meme theory contributes much to the project of explaining the directional change. I discuss each of these points in turn.

17.2.1 The Assumptions of the Tennis-Serve Example

Two assumptions are crucial to the tennis-serve example. The first is that in any population of serves, there is always a unique most attractive serve, in the sense that everyone always chooses to copy it rather than any other serve in the population. All other serves are not copied at all. In other words, perfect selection is assumed.

The second is that variations from the model serve are always small. Even in the version of the example in which people are able to envision what a better serve would be, they can discover only small improvements since “foresight does not extend to the optimal serve itself, merely to slight modifications . . . in that direction” (Gil-White, vol. 2, ch. 16.5.2, p. 326).

I consider the assumption of perfect selection first. It should be obvious how the tennis-serve example depends on this assumption. The basic argument is that despite the fact that every attempt to copy a target serve is imperfect, the population’s mean serve will gradually improve. Why will it improve? As long as the copying errors sometimes produce improvements, there will at some point come into existence a new serve better than the original target serve. Thus, if, at every point in time, everyone is guaranteed to copy the then-optimal serve, the mean serve will improve (Gil-White, vol. 2, ch. 16, p. 324ff). Without the assumption that everyone will copy the optimal serve, however, it doesn’t follow from the fact that copying errors sometimes produce improvements that the mean serve will gradually improve.

It is uncontroversial that a high rate of copying errors can result in cumulative directional change if selection pressures are strong enough. Gil-White (ch. 16, p. 320) cites G. C. Williams’s definition of an “evolutionary gene” as “any hereditary information for which there is a favorable or unfavorable selection bias equal to several or many times its rate of endogenous change” (Williams, 1966), which implies that the acceptable rate of mutation is proportional to the strength of selection. So it should come as

no surprise that if we assume the strongest possible selective pressure—only the best serve gets copied—a very high rate of error is consistent with cumulative directional change. After all, the various copying errors in directions other than the direction of optimal fitness (that is, optimal attractiveness) have no impact on future generations if only the best serve in each generation is copied.

In order for it to be plausible that even given an extremely high error rate, directional change will still occur in a wide range of cases, it must be plausible that selectional pressures will be strong enough. Since it is uncontroversial that sufficiently strong selection can allow for cumulative directional change even if the mutation rate is extremely high (at least assuming that the mutations are relatively small; see following discussion), and since the argument depends on assuming perfect selection, it becomes crucial to see just how problematic the assumption is. I will use the tennis example again to make the point.

Two ways of ranking tennis serves should be distinguished. First, serves can be ranked in terms of their attractiveness, where a serve is more attractive to the extent that people copy it in preference to other serves. Second, serves can be ranked in terms of their effectiveness—their contribution to winning tennis matches.

The existence of such an independent, objective standard of merit for tennis serves perhaps helps to lend surface plausibility to the idea that there is always a single serve that is most attractive to everyone; without an independent standard of merit that everyone is trying to meet, it is much less plausible that all people will always find one candidate meme (“allele”) more attractive than all the other competitors in the population. The case of tennis is special; for many memes it is not the case that there is anything close to a unique standard of evaluation. Moreover, to the extent that selection is strong because everyone is trying to satisfy a unique standard, such as winning tennis matches, the explanation of the strong selection depends on everyone’s sharing a goal, and meme theory thus does less explanatory work.

Yet even in the case of tennis serves, where there arguably is a clear, independent standard of evaluation,³ it is not plausible that there is a unique serve that is most attractive to everyone. First, given the great differences in players’ size and ability, different serves are more effective for different

3. Actually, it is not true that there is an unambiguous effectiveness standard for evaluating tennis serves. What if one serve is more likely to win a point, but is also more likely to injure your shoulder?

players. The serve that is best for a short, powerful man may be different from the serve that is best for a tall, slim woman. An excellent player with strong ground strokes and a weak net game may decide that it would be counterproductive to copy the big serve of someone with a serve-and-volley game.

Second, in playing tennis and choosing a tennis serve, people have motivations other than winning matches. They play to get exercise and for social reasons; they want to avoid injury, impress others, and so on. So even if there were a unique most effective serve (and everyone knew which it was), it would not be true that everyone would try to copy that serve.

Third, an implicit assumption that an effective serve is an attractive serve tends to obscure the point that as a general matter, how attractive a meme is to a person will depend on what other memes the person already has. The way I react to ideas depends on what ideas I am already committed to. This is another reason it cannot be assumed that in any population of competing memes there will tend to be a single meme that is most attractive to everyone.

More generally, which meme among competing candidate memes a person finds attractive may depend on circumstances other than the meme's content. This point illustrates the potential importance of noncontent biases. For example, a serve may be attractive, not because of its effectiveness, but because of the prestige of the person who displays it. We thus have additional factors that may affect the attractiveness of serves; it is an empirical question whether such factors will compete with factors such as effectiveness, thereby decreasing the strength of selection.

I want to emphasize that my concern is not whether tennis serves in particular develop by Darwinian evolution. As noted, whether Darwinian evolution will occur in circumstances where the rate of variation or mutation is extremely high depends on whether selection is strong enough (and whether variations are always relatively small). My present point is that the assumption that selection will be strong enough in a wide range of cases is problematic.

We have seen that even when there is a clear, independent standard of merit for competing memes, we cannot assume that there will be anything close to a unique most attractive meme. Things are presumably much worse when there is no single independent standard or goal that everyone is trying to meet. In such cases, it is often far less plausible that there will be a strong tendency for everyone to copy the same memes; consider men's ties, magazines, religious ideas, desserts, popular music, and so on. At least

on the face of it, what we find is wide variation among different people with respect to what ideas are attractive.

In sum, it is true that given perfect selection, cumulative adaptation can occur even if small errors or variations occur every time an idea is transmitted. But this gives us no reason to be confident that extremely strong selection will be present in a wide range of cases. Moreover, when such selection is present because of humans' pursuit of a shared goal, the explanatory contribution of Darwinian theory is, to that extent, undermined.

I now turn to the second crucial assumption of the tennis-serve example: that variations from the model serve are always small. Without this assumption, there is no reason to think that there will be cumulative change in the direction of the optimal serve. For one thing, if single mutations or innovations can be very large, change need not be cumulative.

It is highly problematic to assume that every variation must be small (relative to the cultural development that is to be explained). The case of tennis serves again illustrates the point. When a weekend tennis player attempts to copy a professional's virtuoso serve, the results are likely to be very far from a good copy. More important, when people try to devise new serves, it cannot be assumed in general that the inventions will be marginal variations on a model. (It is not even true that inventions, in tennis or elsewhere, must be attempts to copy a model.) There is, perhaps, some plausibility to the idea that people won't suddenly come up with a very different tennis serve, but this likely has more to do with the limitations imposed by the human body and the rules of tennis than with a general truth that ideas change in small steps.

In general, ideas can be revolutionary, path-breaking. Scientists, inventors, and artists certainly build on the work of others, but there is no reason to think that large changes in ideas always come about by an accretion of small changes, beginning from a target idea and moving through a chain of descendants. Single developments need not be small. And the process of innovation need not be cumulative. An innovator may be influenced at once by many different ideas from widely disparate sources. These points will be important when I turn in the next section to the explanatory power of a Darwinian theory of cultural change.

17.2.2 What's the Best Explanation?

In the preceding section we saw that Gil-White's argument from the tennis-serve example depends on strong empirical assumptions. I suggested that those assumptions are not plausibly close to the truth in the case of

tennis serves, and are much less plausible for memes in general. In this section I turn to my second, more fundamental point: Even when the necessary conditions are satisfied, so that cumulative change can occur, it does not follow that a Darwinian account is the best explanation of the change.

The tennis-serve example shows that a high error rate need not rule out cumulative directional change, but more is needed to defend meme theory successfully. Even when there is, broadly speaking, Darwinian evolution—differential survival of elements and cumulative change in the direction of complex apparent design—the best explanation of the change may be thinkers' deliberate innovation in pursuit of their conscious goals. If goals do all the work, appeal to Darwin is idle.

Suppose the meme theorist responds that whatever the explanation of the differential transmission of ideas, it is still true that the ideas that are around today are the descendants of the ones that succeeded in spreading the most. (Even this tautological-sounding claim may not be true since it may not be true that current ideas are the descendants, in any relevant sense, of earlier ideas.) According to this response, the fact that in some cases ideas spread because of our goals does not undermine meme theory any more than any other explanation of why some memes have greater fitness than others.

The problem with this imagined response is that if this were all that meme theory claimed—that the memes around today are the ones that spread the most—it would have nothing to offer. Meme theory needs to offer a distinctive explanation of why ideas spread, an explanation that competes with, or contributes something beyond, the goal-based account.

The tennis-serve example again illustrates the point. Rather than selection by blind forces, we have deliberate adoption of the most effective serve. And, in the more realistic version of the example, new serves are generated by self-conscious, successful innovation. Thus, the goal-based account of cultural change is a good candidate to explain the development of tennis serves. It therefore needs to be argued rather than assumed that Darwinian theory is a better explanation, or at least that it has something to add.

Gil-White suggests that the conclusions he draws from the tennis-serve example apply to technology in general, where "design has accumulated gradually," and thus that "selectionist approaches will be significant for historical explanations" (vol. 2, ch. 16, pp. 327–328). It is important to see why more argument is needed here.

First, although there is a weak sense in which technological change is cumulative (new technologies draw on existing knowledge), it is not clear that it is cumulative in the relevant sense (built up by successive additions). There can be sudden large breakthroughs because of one thinker's great idea (or because of an accident), as opposed to the accumulation of small variations through differential reproduction. Although researchers obviously draw on past ideas, each new idea need not be formed by making a change to a single past idea. And one can decide to scrap a whole project or design and begin from scratch.

Second, more importantly, in the case of technology there is a strong *prima facie* case that deliberate innovation is the best explanation of many developments. At least on the face of it, the reason that jet-engine technology has been invented and widely transmitted is that it satisfies human goals. If we contented ourselves with the thought that jet-engine memes are good at reproducing themselves in human minds (or have spread as the result of biases in the copying of ideas), we would miss a fundamental part of the explanation. (It is instructive to compare trying to predict the future course of technological development by considering, on the one hand, what goals humans have and, on the other, which memes are good at getting themselves reproduced by human minds and how human copying of ideas tends to be biased.)

It *could* be true that although people devote great time and energy to research and although they try to adopt the best models that research produces, a goal-based account of technological change is inadequate. How could this be? One way is that it could turn out that despite all efforts at deliberate innovation, the important steps forward tend to be the result of small copying errors—accidental mistakes in transmitting ideas. A different possibility is that even if the best explanation of each individual's decision is in terms of his or her goals, population-level phenomena, such as group selection, are crucial to the explanation of which technologies spread (Boyd & Richerson, 1985, 1989, 2000). A third possibility, which I take up in the next section, is that ideas could spread not because they serve our purposes, but because they are "catchy" (in a precise sense that I will explain). Again, it does not follow from the existence of cumulative change that Darwinian theory adds anything to the explanation of that change.

17.3 The Meme's-Eye Perspective

I now turn to the relevance of the "meme's-eye perspective" to the question of explanation. The idea is that the "interests" of memes explain the

direction of cultural evolution—the analogue of the selfish-gene idea in biology. The important point here is that the selfish-meme approach⁴ is a potentially important source of explanatory power for a Darwinian theory of culture; in particular, it is a main way in which meme theory may challenge goal-based explanation of cultural change.

I am not suggesting that Darwinian explanation must be adaptationist. As we will see, another possibility is that certain ideas spread because of social learning biases, for example, a human tendency to copy prestigious people. The point is rather that the explanatory payoff of Darwinian theory, whether adaptationist or not, cannot be taken for granted.

We have seen that meme theory needs to offer distinctive explanations for the spread of memes. To the extent that the fundamental explanation of the creation and transmission of memes is our pursuit of our goals, meme theory is doing no work. In the case of biology, Darwinian theory is an answer to the grand puzzle of how complex, apparent design has come about without any intelligence or design at all—a question that does not arise in the case of culture. Darwinian theory also provides answers to smaller-scale questions of why biological evolution takes particular directions. In the case of culture, there are analogous smaller-scale questions. For example, why does a particular idea take a long time to come about and then develop independently in many different places? Why do some ideas spread rapidly and persist for centuries? It is such questions of why cultural evolution proceeds in particular directions and at particular rates that meme theory must purport to answer.

In neo-Darwinian theory, one highly influential approach, made famous by Richard Dawkins (1976/1989), is to explain the course of biological evolution by taking the perspective of the gene, rather than that of the organism, group, or species. Roughly speaking, according to this “gene’s-eye perspective,” characteristics have evolved, not because they are in organisms’ or species’ interests, but because the genes for those characteristics were more successful than other genes at reproducing themselves in the local environment.

When Dawkins introduced the notion of the meme, his suggestion was that a meme’s-eye perspective could offer explanatory power analogous to that of the gene’s-eye perspective. The idea is that some features of culture are better explained by the memes’ effectiveness at reproducing themselves than by what human creators of culture want or intend or have as goals. For example, we might better explain why some fad or prejudice or tune

4. I use “selfish-meme approach” and “meme’s-eye perspective” interchangeably.

content biases are potentially an important way to challenge goal-based explanations of cultural evolution.

Noncontent biases are explanations of the spread of ideas that do not appeal to the ideas' content. There is, however, an important distinction within the class of content-based explanations. On the one hand, as I have emphasized, ideas that are true or useful to humans can be spread because of humans' conscious goals.⁵ On the other, ideas can spread because their content takes advantage of other features of human psychology—despite or regardless of humans' conscious goals or what is good for humans. That is, ideas may be catchy or infectious to human minds or brains.

We have to be careful to distinguish two senses of "catchy." In a broad sense, an idea is catchy if it has a tendency to spread. In this sense, the fact that a meme spreads as a result of its catchiness (as opposed to spreading as a result of noncontent biases) does not settle the question of the best explanation of the meme's spread. It might be that the best explanation of the meme's spread is our deliberate pursuit of our goals. Call this broad sense of "catchy" the *vacuous* sense.

In a narrow sense of "catchy"—the *interesting* sense—an idea is catchy if the idea's content makes it good at reproducing itself because of features of human psychology other than humans' deliberate pursuit of their goals. So advertising jingles, clichés, and religious cults are catchy in the interesting sense. In contrast, when people deliberately invent and spread an innovation—for example, the electric light—because it serves their interests, the explanation of why the innovation spreads is not its catchiness in the interesting sense, but our skillful and conscious pursuit of our interests. (In the vacuous sense of catchiness, the electric-light idea is catchy. But in this sense, catchiness is not the explanation of why the idea spreads. So notice that in neither the vacuous nor the interesting sense is catchiness the explanation of why the electric light spreads.)

Once we make the distinction between vacuous and interesting catchiness, we can see that the notion of interesting catchiness provides an important way in which Darwinian meme theory can compete with goal-based explanations of cultural evolution. Meme theory can try to show that the best explanation of why some ideas spread is not our conscious

5. It is an oversimplification to say that goal-based explanations must be content-based. Although goal-based explanations typically depend on features of ideas' contents, such as their truth or usefulness, goal-based explanations need not do so; in some circumstances, it will serve one's goals to adopt ideas for reasons other than their content.

spreads by appealing to its “catchiness” than by appealing to human beings’ intentions, goals, interests, and welfare. Just as a highly infectious virus may spread because it is good at taking advantage of features of our physiology rather than because we deliberately spread the virus to promote our goals, a highly infectious idea may spread because it is good at taking advantage of features of our psychology (other than our goals), rather than because we deliberately spread the idea to promote our goals.

Dawkins (1976/1989, pp. 196–200) suggested that since the memes that are good at reproducing themselves will tend to be the ones that are around today, we can think metaphorically of the memes as “trying” to reproduce themselves—as having an “interest” in doing so. Thus, the meme’s-eye perspective explains the spread of ideas in terms of the interests of memes rather than of humans. Furthermore, meme theory might try to turn the tables on goal-based accounts, explaining why we have come to have our goals and intentions in terms of the interests of memes. (Of course, even if it is true that the meme’s-eye perspective explains why we have certain goals, it doesn’t follow that the goals are not what is now driving cultural change. Compare: biological evolution explains how memes came to be in the first place; it doesn’t follow that memes are not now important in explaining cultural change, and possibly even in explaining which genes are selected for.)

It is dangerous to ignore the issue of meme theory’s explanatory power. For example, Gil-White criticizes the meme’s-eye perspective by emphasizing the importance of *noncontent biases*: whether and how frequently a meme gets transmitted can depend on features other than its content. It may be, as mentioned, that humans tend to adopt ideas that derive from prestigious members of the community. Or, how likely humans are to adopt an idea may depend on the frequency of the idea and of other ideas in the population. (As indicated earlier, such noncontent biases may conflict with the assumption that everyone will try to copy the most effective tennis serve.)

Why should the idea that memes can spread not just because of their content but also because of noncontent biases count as a criticism of the meme’s-eye perspective? This would be an effective criticism if the gist of the perspective were the claim that only a meme’s content is relevant to its spread. As I will elaborate, however, much of the importance of the perspective lies in its potential to challenge a goal-based account of the spread of ideas. Indeed, the significance of noncontent biases themselves cannot be understood without asking what explanatory payoff Darwinian theory provides over a goal-based account. Like the meme’s-eye perspective, non-

content biases are potentially an important way to challenge goal-based explanations of cultural evolution.

Noncontent biases are explanations of the spread of ideas that do not appeal to the ideas' content. There is, however, an important distinction within the class of content-based explanations. On the one hand, as I have emphasized, ideas that are true or useful to humans can be spread because of humans' conscious goals.⁵ On the other, ideas can spread because their content takes advantage of other features of human psychology—despite or regardless of humans' conscious goals or what is good for humans. That is, ideas may be catchy or infectious to human minds or brains.

We have to be careful to distinguish two senses of "catchy." In a broad sense, an idea is catchy if it has a tendency to spread. In this sense, the fact that a meme spreads as a result of its catchiness (as opposed to spreading as a result of noncontent biases) does not settle the question of the best explanation of the meme's spread. It might be that the best explanation of the meme's spread is our deliberate pursuit of our goals. Call this broad sense of "catchy" the *vacuous* sense.

In a narrow sense of "catchy"—the *interesting* sense—an idea is catchy if the idea's content makes it good at reproducing itself because of features of human psychology other than humans' deliberate pursuit of their goals. So advertising jingles, clichés, and religious cults are catchy in the interesting sense. In contrast, when people deliberately invent and spread an innovation—for example, the electric light—because it serves their interests, the explanation of why the innovation spreads is not its catchiness in the interesting sense, but our skillful and conscious pursuit of our interests. (In the vacuous sense of catchiness, the electric-light idea is catchy. But in this sense, catchiness is not the explanation of why the idea spreads. So notice that in neither the vacuous nor the interesting sense is catchiness the explanation of why the electric light spreads.)

Once we make the distinction between vacuous and interesting catchiness, we can see that the notion of interesting catchiness provides an important way in which Darwinian meme theory can compete with goal-based explanations of cultural evolution. Meme theory can try to show that the best explanation of why some ideas spread is not our conscious

5. It is an oversimplification to say that goal-based explanations must be content-based. Although goal-based explanations typically depend on features of ideas' contents, such as their truth or usefulness, goal-based explanations need not do so; in some circumstances, it will serve one's goals to adopt ideas for reasons other than their content.

goals, but the ideas' exploitation of other features of our psychology. Meme theory can even try to argue that our goals are themselves the product of our previous infection by catchy memes. The meme's-eye perspective is precisely the attempt to show that it is the meme's interests rather than ours that are in the driver's seat.

If we don't recognize the challenge that a goal-based explanation presents to other accounts of cultural development, we will not appreciate the importance of the meme's-eye perspective (as Gil-White's position illustrates). For the meme's-eye perspective, far from merely insisting on content-based explanations of the transmission of ideas, offers a way of challenging goal-based explanation, a central case of content-based explanation. Similarly, the importance of noncontent biases is better brought out by seeing them as another way of challenging goal-based explanations than by seeing them as a challenge to the meme's-eye perspective. Indeed, the reason it is appropriate to talk of *biases* in the transmission of ideas, as opposed to, say *methods, strategies, or policies*, is that it is implicit that an explanation in terms of biases is supposed to compete with an explanation in terms of deliberate choices. From this point of view, noncontent biases are an important ally, rather than a refutation, of the meme's-eye perspective.

17.4 Conclusion

Meme theory has the potential to challenge or complement other explanations of human cultural evolution, but it is crucial for theorists not to take the theory's explanatory power for granted. In order to defend meme theory successfully against the charge that the mutation rate is too high, it is not enough to show that if we assume perfect selection and relatively small variations, cumulative directional change can still occur. These assumptions themselves must be supported, and, more important, so must the substantive claim that Darwinian theory explains cultural change better than the commonsense goal-based account.

Noncontent-based explanations of cultural evolution are indeed important. The point of the meme's-eye perspective, however, is not to insist that explanations of cultural evolution must appeal only to the contents of ideas, but to oppose explanations in terms of our goals. Thus, noncontent-based explanations do not undermine the meme's-eye perspective. The meme's-eye perspective offers a possibly fruitful way for Darwinian theory to challenge the commonsense explanation of cultural change in terms of our skillful pursuit of conscious goals. In fact, the importance of non-

content-based explanations is better understood once we see them as another, complementary way of challenging goal-based explanations of culture.⁶

Acknowledgments

I am grateful to Andrea Ashworth, Francisco Gil-White, and Paul Seabright for helpful comments. I owe special thanks to Susan Hurley and Harry Litman for detailed suggestions and discussion.

6. See comments on this chapter by Blackmore (vol. 2, ch. 19.12, p. 406) and relevant discussion by Chater (vol. 2, ch. 18, p. 355). ED.