Crowding out Memetic Explanation

Rosa Cao*†

Memes have been proposed to explain wide swathes of human culture and language use. I argue that what is really doing the explanatory work in many of these cases is a basic mechanism of information transmission, which is distinct from memetic evolution by natural selection in significant ways. Perhaps the most significant of these is that information transmission depends primarily on the interests of the users of information, rather than the reproductive interests of the informational entities—'memes'—themselves. Although my main target is memetic approaches, this argument also applies to some other, nonmemetic, theories of cultural evolution.

1. Introduction. "Memes are *ways*," Dennett (2017, 206) tells us: ways of behaving, doing, or making. They include paradigmatically human tools such as spears and words, cultural items such as songs and internet images, as well as the skills and strategies of other creatures. We can immediately see then that memes are intimately bound up with how agent-like things go about their business of behaving, doing, and making.

This should be no surprise—after all, memes are informational, not in some kind of dry formal sense but in a richly semantic way. If memes were nothing more than a kind of statistical or Kolmogorov information, then everything and nothing would be a meme, solely in virtue of physical structure.[1] That memes are more than this is what makes them interesting and plausible candidates for cultural explanation. They already have meaning,

*To contact the author, please write to: Department of Philosophy, Stanford University, Stanford, CA 94305; e-mail: rosacao@stanford.edu.

†My thanks to Daniel Dennett, Patrick Forber, Peter Godfrey-Smith, David Haig, Enoch Lambert, Ron Planer, Charles Rathkopf, Gill Shen, Shawn Simpson, Jared Warren, and audiences at Stanford and University of California, Merced, for helpful discussions and feedback.

1. One might think that this is the most charitable interpretation of declamations such as Wheeler's "it from bit" (1989).

Philosophy of Science, 87 (December 2020) pp. 1160–1171. 0031-8248/2020/8705-0033$10.00
Copyright 2020 by the Philosophy of Science Association. All rights reserved.

https://doi.org/10.1086/710518 Published online by Cambridge University Press

and so it is no mystery how they can explain meaningful practices that rely on semantic content. However, this also means that their very existence depends on the kinds of information users who deal in semantic information themselves.

What is distinctive about true memetic explanations is that they are evolutionary in a very direct way: the memes themselves are units of selection, and what is explained is how those units of selection are distributed, as well as how they arose through an entirely orthodox process of evolution by natural selection (ENS). Inasmuch as discussions of cultural evolution that do not directly appeal to memes (e.g., Henrich 2015) share this feature—of locating agency and fitness with the elements of culture themselves rather than the human populations in which they reside—the considerations presented below will also apply to them.

There are different formulations of ENS, and many earlier challenges to memetic explanation focused on the Dawkinsian framework (Dawkins 1976) of replicators and interactors.[2] I will use the less restrictive formulation due to Lewontin (1985), with the further dimensions of flexibility explored by Godfrey-Smith (2009). This formulation has three main components, each of which may be realized by many mechanisms, and each of which may occur to varying degrees: variation, heritability, and differential fitness. Depending on how well a process in the world exhibits these three characteristics, it will be more or less accurately described by ENS, and thus ENS will offer a better or worse explanation of the outcomes and trajectory of that process. Another way of putting it: if we are interested in populations of entities and how they change over time in conjunction with their environments, ENS will explain just that part of the change over time that is due to the interaction of these three components, while leaving the rest to be accounted for by other factors.

Since memes are to be the units of selection, to the degree that an agential or teleological perspective toward the entities involved in evolutionary processes is appropriate, we will in fairness adopt them toward memes and hence talk of the "selfish meme," their evolutionary and reproductive goals, and so on.

So here is a recipe for a memetic explanation. First, identify the explanatory target of interest. This might be a distribution explanation in which we want to explain changes in the distribution of a word, behavior, or artifact over time. Or it might be an origin explanation in which perhaps we want to explain the (evolutionary) emergence over many generations of an interestingly structured or complex informational entity. Then, for the explanation to go by way of ENS, we need clearly identified units of selection (i.e., clear

2. Sperber (2000), e.g., argued that memes were not replicators because their reproduction was insufficiently faithful. Also see Sperber and Claidière (2006).

https://doi.org/10.1086/710518 Published online by Cambridge University Press

individuation conditions for the meme in question) and mechanisms of inheritance, variation, and selection. I focus on the issue of inheritance, since it is in that context that memetic reproduction seems most easily confused with informational transmission, at the cost of providing pseudoexplanations.

It is important to note in this explanatory context that reproduction should be either the only or the primary means by which memes reoccur in the target phenomenon. This is because if genuine memetic reproduction accounted for only a small proportion of the occurrences of a meme, then it would do a bad job of providing distributional explanations—it would account for less and less of the distribution of features that is actually observed in the population over time.

For comparison, imagine what would happen to our evolutionary explanations if there were ways for biological entities to recur without biological reproduction. If creatures with certain traits just randomly popped up in significant numbers, that would wreak havoc with the predictions provided by ENS of the distribution and origin of those traits, and the explanatory import of ENS in the explanation would be diluted, if not overwhelmed altogether. Perhaps something similar to this happens when there is significant migration into and out of the population of interest—then it seems that we do not have good boundaries on the relevant population.

2. The Individuation of Memes. How do we individuate memes, for the purposes of formulating our explananda? And how do we count them in the process of providing the explanation? Genes are difficult enough to keep track of—and we know what they look like and what they are made of biochemically.

One immediate difficulty arises from memes' informational nature—it is paradigmatic of informational entities that they are substrate independent. For example, Dennett (2017) tells us that words are paradigm memes and that manifestations of them that do not share common physical features nonetheless count as tokens of the same type: shouted, whispered, written, or stored away in memory for future use—it makes no difference.

I can see only one way to make sense of this fact, and that is by realizing that what makes these tokens the same is simply that they are all *taken* to be of the same type, by some interpretive mechanism belonging to an information user. It is only from the perspective of that user or that mechanism that things with potentially quite different physical characteristics could count as the same meme, with the same informational content.

Moreover, the more physically different manifestations of the same meme are allowed to be, the more sophisticated that interpretive mechanism must be. If the interpreter did not have an auditory system as well as a visual system and the ability to read, then a spoken word and one written down would not be the same meme and would not have the same effects. For many of the

https://doi.org/10.1086/710518 Published online by Cambridge University Press

most interesting memes, that mechanism is going to be something like an agent with multiple sensory-perceptual capacities and something like a mind of its own.

This might seem troubling. Are we introducing an explanatory circle involving agents, representations, and goals every time we invoke information use? On my account of information use, these notions of agency, information, and goals or interests do necessarily come and go together. And so appeals to semantic information do not automatically provide a reductive explanation of behavior, although they can. That these concepts all hang together need not detract from the usefulness of informational explanations more generally however.[3]

Notice that the bare idea of "taking as" is a functional notion that, on its own, does not require much cognitive sophistication on the part of the information user. What does require sophistication is the capacity to take physically diverse things as functionally the same, by expressing similar response profiles to very different physical manifestations. That requires the ability to marshal different sensory-perceptual capacities and coordinate their deliverances in a unified way. Thus what binds the physically diverse instances of a meme together is that they play the same functional role for one or more receivers (and only when this is the case will the theorist find some additional explanatory payoff or convenience in bundling together diverse physical entities as instances of a single meme).

So the kind of information we are talking about when we talk about memes as informational entities is necessarily information from some perspective. Information is necessarily information *for*. And that means that wherever we see memes, other agents are necessarily lurking nearby, and we will have to show that their intentions are not ultimately what explain the target phenomenon.

I will come back to this shortly, but for now, the key point is just that there are many ways to individuate memes, depending on which perspective we take. And so a scrupulous memetic explanation must be careful to rely on memes that are picked out with respect to the very same mechanisms.

3. Do Memes Really Explain Anything (New)? Memes need not resemble each other physically, so long as some suitable agent takes them to be the same, informationally. But if memetic explanations are to be genuinely Darwinian explanations, the appearance of new instances of memes must be the consequence of reproduction by earlier instances. That is, the number of instances of a type must depend on the features of earlier instances of that

3. I have already assumed (as is common in this literature) a reductive notion of goals and interests given by the usual evolutionary teleonomy framework. Agents, too, can be minimally defined as bearers of these simple interests, and so on, in the style of Dennett's homuncular functionalism (see Dennett 1978; Cao 2012).

https://doi.org/10.1086/710518 Published online by Cambridge University Press

type. Anything like spontaneous generation must be ruled out, or else this distinctive fitness relationship at the heart of ENS will be lost.

While there are no cases of spontaneous generation in biology (except the origin of life), the same is not true of memes. There does not seem to be a single memetic tree through which all memes may trace their reproductive lineage. Some critics have argued that memes rarely if ever reproduce in the required ways for ENS to be a good account of their distribution (Sperber 1996, 2000; Godfrey-Smith 2009). Rather, they occur, and they reoccur, but the mechanisms responsible are not best described as reproductive.

As Sperber (2000) recounts, one of the ways in which a behavior (e.g., laughter) may recur is through a preexisting motor program inherited by the usual biological channels. Another way is by the expression of a previously acquired concept, such as "five-sided star." We may be triggered to produce a star by copying it, but it seems that we are not copying it in the same way that we would copy an unfamiliar figure. For one thing, it is fast, automatic, and much lower error. Sperber acknowledges that many behaviors will be the result of complex combinations of genuine copying and triggering effects on prior knowledge and inferential capacities, so there is a spectrum here. But pure copying is, he argues, rare.

It certainly seems as if informational entities occur and reoccur. Figure 1 shows three of the ways this can happen. The question is whether their patterns of occurrence and the features they exhibit are best explained by ENS, as illustrated in the best case scenario (fig. 1*b*) or whether instead they are better captured by a traditional picture appealing to the information-using capacities of agents sending and receiving information for their own purposes (fig. 1*a*). A third, Sperberian, alternative to both of these possibilities explains recurrence by appealing to the constructive and creative capacities of hosts satisfying independent constraints not determined by the original meme, which functions as a mere trigger (fig. 1*c*). In both cases 1*a* and 1*c*, memetic explanation is unappealing. For 1*a*, it seems otiose—we can do just as well by appealing to the interests of the agents involved in sending and receiving informational signals. Perhaps an example here would be the relaying of alarm calls. Notice that when we talk about the interests of information-using agents, we include not just the explicit, person-level interests of humans but also their evolutionary interests, whether or not those register at a conscious level or align with explicitly held goals.

In figure 1*c*, memetic explanation seems inadequate to explain the creation of a new informational entity. As Sperber and others have elaborated at length, there may be many parents, or none, and there is no guarantee of higher similarity between triggering signals and the creations they trigger than between that creation and other triggers with different informational features (Sperber 1996; Fracchia and Lewontin 1999; Wimsatt 1999). Recreations may be the result of inference, or construction, rather than any kind

https://doi.org/10.1086/710518 Published online by Cambridge University Press

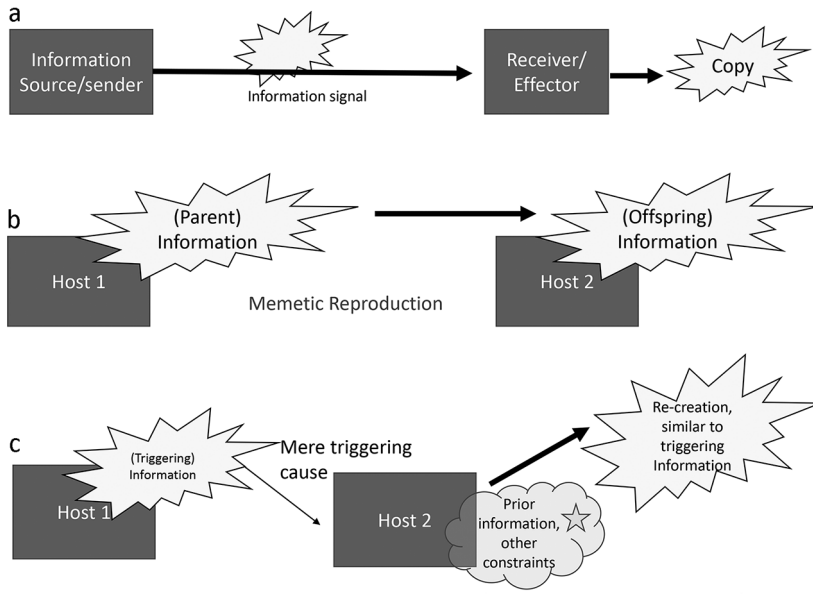


Figure 1. Three ways in which informational entities can recur. Only *b* is reproduction of the kind required for evolution by natural selection. *a*, Standard information transmission explanation: agency and causal responsibility rests with the users of information, not the information itself. *b*, Genuine memetic reproduction: hosts serve merely as environmental scaffolding for memetic reproduction, and the causal arrow goes from informational meme parent to meme offspring. The locus of agency shifts to the informational signal itself. *c*, Sperberian re-creation: the putative meme serves as the mere triggering cause for a host to re-create a new instance, on the basis of resources it already possessed. Primary causal responsibility rests with the host, not the putative meme. Color version available as an online enhancement.

of copying that depends for its specific structure on the specific structure of the trigger.

Information use and Sperberian explanations also differ from one another in their emphasis. For Sperber, the agent's background knowledge and inferential and other cognitive capacities are important as the best explanation for the creation of new informational tokens, but it may still be the case that the re-creation of these new instances does not particularly benefit the agent or host. For the information-transmission explanations, the agent or effector is important not so much for what it knows but for the fact that it is the user of the information—the one who benefits from its propagation. It is only with respect to the user's capacities, its perspectives, decoding abilities, and purposes, that there is information in the signal at all.

https://doi.org/10.1086/710518 Published online by Cambridge University Press

So the two bad cases for memetic explanation illustrate two worries. First, it may be that the best attribution of agency is to the users of information and not the informational entities themselves. This includes as-if agency in Dennett's sense—which may be aptly attributed to simple information users without rich psychological lives. And second, it may be that causal responsibility for the generation of a new informational entity rests with the hosts and other constraints they may be satisfying or information that they had independently stored. In that case, the triggering "meme" would be merely a permissive or triggering cause and not a specific (informational) cause of the informational shape of the outcome.

Meme enthusiasts might argue that regular information transmission does not crowd out memetic explanation but is rather an instance of it. To this I would respond: the attribution of agency, including as-if agency, is competitive. There is a certain "amount" of agency that will explain some phenomenon, and if we assign it at one level, then that will obviate the need to assign it elsewhere as well. Why posit two sets of agents when one will do?

Thus, if the decisions of a group are fully accounted for by the decisions of the individuals within the group acting as agents, then there is no need to also posit group agency. Similarly, if a perfectly good explanation can be given of my reproductive activities in terms of my biological interests, then it seems that it would be otiose—within the same explanation, targeting the same phenomenon—to also assign agency to my cells.[4] This is not, of course, to insist that equally good explanations could not in some cases be given from either perspective—it is just that they are in some sense equivalent, and we do not need both at the same time.

The selfish gene was threatening because it seemed to undermine our own agency, interests, and intentions in favor of those of our genes. Conversely, when the intentions and goals of agents already explain why they respond to and produce the informational signals that they do, then it seems to complicate things unnecessarily to ascribe agency to those signals themselves. We want the intentions of the memetic host to be background conditions, scaffolding the reproduction of memes; we do not want them to be directly operative as the core of the explanation.

Let me sum up the discussion so far. There are (at least) two dimensions along which a putative memetic explanation may be more or less apt. The first is where we place the locus of agency and the evolutionary interests that are being advanced by the multiplication of the informational entity. For a satisfying memetic explanation, these should lie primarily with the meme itself, not with the hosts or users of information.

4. For similar arguments in different contexts, see Godfrey-Smith's (2009) discussion of the de-Darwinization of collectives, as well as Kim's (1993) causal exclusion argument against nonreductive materialism.

https://doi.org/10.1086/710518 Published online by Cambridge University Press

The second dimension is where to place the responsibility for the creation of new instances of a putative meme. In some relevant sense, we want the information content of the parent to determine the information content of the offspring, and moreover, we need that information content to be the same between parents and offspring (or at least more similar than between unrelated members of the population). This happens when the host serves as a kind of unwitting scaffolding for the reproduction of the informational entity. The host cannot become too engaged as to make too many of its own modifications to the re-creation, nor can it be so sloppy as to lose too many features of the original. It cannot be inspired by informationally irrelevant causes, nor can it rely primarily on its own internal resources to generate the particular features. A fine balance must be maintained in which the host makes enough of its resources available for the reproduction of the meme but not so many that the best explanation for the features of the re-creation depends on the specific features of the host, with the putative parent merely playing a triggering or permissive role.

A comparison might be made here with parity arguments concerning genetic versus environmental contributions to biological reproduction, many of which are collected in Oyama, Griffiths, and Gray (2001). Developmental systems theorists argued that genes are not a distinctive or privileged cause of phenotypic outcomes, or heredity, but rather that all parts of the developmental system contribute to outcomes in a concerted and interactive way. On their view, the "informational" gene is on a par with respect to its explanatory and causal responsibility with other parts of the cell and environment. In that debate, one set of responses appealed to the specificity and systematicity of genetic effects to break the putative parity between genes and environment (see, e.g., Waters 2007).

A version of my worry for informational explanation more generally (and memetic explanation in particular) is that we are putting the same undue focus on certain informational contributions, at the cost of ignoring other causes with genuine explanatory responsibility of their own. In Sperber's examples, some of the other causes he points to may also be informational but reside within the cognitive system of an agent, and their contribution is not merely to be a template for copying.

Can the same response as the one given on behalf of treating genes as privileged causes in some way be given here? I think not, because, if anything, the same arguments militate toward treating the biological agents as privileged causes, rather than the memes whose reproduction they scaffold. It is the biological agents that are systematically difference makers across a range of outcomes, and it is the biological agents the tweaking of whose cognitive systems can create very specific changes in the informational outputs produced. If we are scaffolds for memetic reproduction, we are scaffolds that may turn on the construct being built at any time, changing them

https://doi.org/10.1086/710518 Published online by Cambridge University Press

to our own purposes while only rarely merely being used for theirs. If the environments in which genes replicate were to have the same kind of agential powers, what we would see is a kind of ubiquitous threat of genetic engineering and intervention. Any difference between fitness and "genes that the agents aim to produce" would collapse; in such a world, dominated by intentional and direct genetic manipulation, ENS, too, would seem otiose.

4. The Rarity of a Useful Memetic Explanation. Memetic explanations live in a crowded space, competing with non-Darwinian informational explanations on one side and Darwinian explanations in terms of populations of information users (and not memes themselves) on the other. There is also an undermining worry present for all three kinds of explanations that the putative central player—the information—is not actually doing much work, and we are misattributing responsibility to information transmission or reproduction when instead a sufficient explanation is already provided by attending to the causal or mechanistic scaffolding for the process.

The space of explanations is illustrated in figure 2, with the locus of ascribed agency on the vertical axis and the degree of causal responsibility

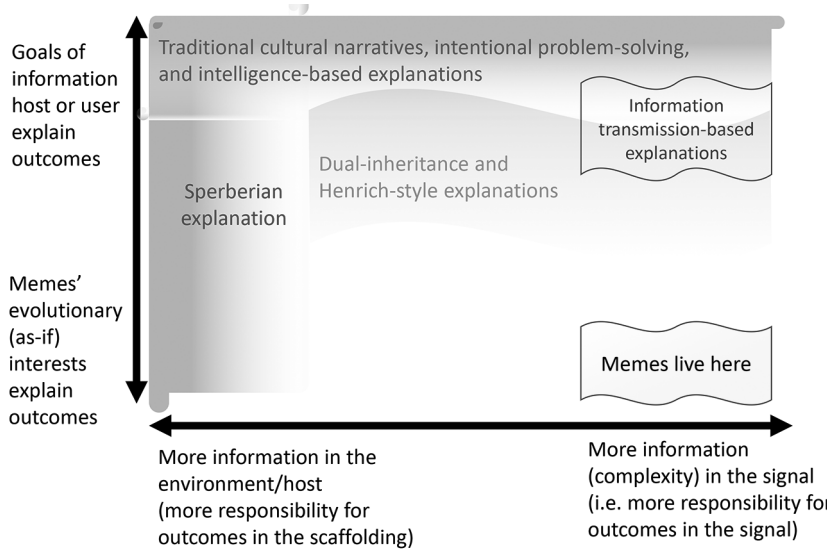


Figure 2. Pulling away from the memetic corner. Note that for Boyd/Richerson/Henrich-style explanations (*upper middle*), it is still the evolutionary fitness of the information users that primarily explains the outcomes of cultural evolution, not the as-if interests of the cultural variants themselves. Color version available as an online enhancement.

https://doi.org/10.1086/710518 Published online by Cambridge University Press

attributed to informational entities as opposed to their hosts or environments on the horizontal axis. While the horizontal and vertical dimensions are logically independent, it may be that as a matter of fact, they go together in such a way as to pull phenomena away from the memetic corner. Why? Because of the earlier observation that in a world with no agents, and no users of information, there are also no memes. So we need somewhat sophisticated agents in order to make sense of informational entities in the first place.

At the same time, it has to be that one cannot offer an equally good explanation just by appealing to the goals and psychology and agency of the users. Memetic explanation will be superfluous if the agency of information users exhausts the phenomenon. So the remaining candidates for memetic explanation are only those cases in which there is enough agency around that there are clear memes, but that agency is directed elsewhere—toward some goal other than the advanced by the expression or recurrence of the memes—and the rational behavior of the agent(s) relative to which the memes are individuated cannot explain those memes' distribution or origin.

Why should the (as-if) agency of information users be explanatorily prioritized over the (as-if) agency of memes, if both are equally capable of explaining the phenomena? One thought is that perhaps memes are simpler and so provide a kind of reductive appeal. I think any reduction of complex cognitive phenomena to informational memes is an illusion, because of the earlier point about individuation: memes are only visible from the point of view of something that is capable of taking them as having a particular informational content, so our explanation will still rely on taking some significant cognitive capacities for granted.

Another thought is that memes may be no better at explaining some phenomena than a direct appeal to the interests of an information-using agent, but perhaps they are no worse. Here again, I think the point fails because biological agents are in a sense better entities. They are more inferentially deep, in the sense that they are visible from more points of view. All else equal, we should prefer to appeal in our explanations to preexisting entities that have already proved their worth, rather than posit fragile new ones that are only visible from quite particular points of view in quite particular explanatory contexts.

There is another pivot available to the meme advocate. And that is to reclaim the physical nature of memes. Sterelny (2006, 155), in making this point about artifacts, skills, and other technologies says, "we have to abandon the idea that memes are *ideas*." This point is quite general: some memes are adopted because of their physical features—ways that they tickle us that have nothing to do with our information-processing and agential capacities. Some items (whether they be alarm calls, words, or pictures) appeal and propagate in virtue of particular physical features that they have, so that even if the

https://doi.org/10.1086/710518 Published online by Cambridge University Press

relevant host population could recognize the item as being informationally the same despite having different physical features, their appeal and thus fitness would differ. A verbal description of a furry creature with a large head and big dark eyes does not have the same visceral effect or multiplicative potential as the picture of a baby seal gazing imploringly at you. Because our intentional, information-processing agency is less involved in these cases, more room is left for memetic explanation.

The space of explanations shown in figure 2 also allows us to situate approaches to cultural evolution that do not appeal explicitly to memes but that nonetheless rely on a similar thought: that somehow we are not responsible for our "clever" solutions to a wide range of problems. For example, Henrich (2015) argues that it is cultural evolutionary processes that explain many impressive human achievements, rather than the intellectual problem-solving capacities of human individuals themselves.

This message can and should be separated from the claim that cultural evolutionary processes are acting on the fitness not of humans but instead of their cultural artifacts, in whatever form. Nongenetic inheritance within human populations is one thing; cultural items with their own evolutionary interests is quite another.

So the kinds of explanations advocated by those such as Henrich, Boyd, and Richerson (e.g., Boyd and Richerson 1985; Richerson and Boyd 2005) should sit in an intermediate position on the vertical axis in figure 2. We are not all the way at the top, because the self-aware intentions and explicit understandings of these creatures do not directly explain the phenomena. Rather, much of the responsibility for the differential uptake of different cultural variants lies with psychological and physiological quirks and capacities that operate semi-independently from the conscious intentions and explicit interests of their owners. Crucially however, the bearers of fitness—the (as-if) agents with evolutionary interests of their own—are biological organisms, not their artifacts or cultural variants. There are many reasons to be reluctant to call those artifacts or cultural variants memes; for the considerations surveyed here, we should consider them adaptations found in biological populations, rather than forming Darwinian populations of their own.

I conclude by saying something about when it is that memetic explanations are apt. Dawkins (1993) suggests that they are most useful when they provide evolutionary explanations of apparent host-level irrationality where none was available before. To this I would add an observation from Ullman-Margalit on hidden-hand explanations: that their explanatory interest increases "when there is a difference in *type* between the overall pattern to be explained and the individual actions which are supposed to bring it about, as well as with the complexity of the intermediary process" (1978, 267). Ullman-Margalit describes those as explanations that turn on the unintended consequences of human agency, cases in which "numerous individuals each

https://doi.org/10.1086/710518 Published online by Cambridge University Press

busily doing his or her own private narrow bit, such that an overall design, unsought as well as unforeseen by them, is seen to emerge" (271).

Such cases are plausibly rare but striking when they happen—so striking that our saliency heuristic may cause us to overestimate their prevalence. I suspect that such striking cases are the exception and that, for the most part, the presence of intelligent and as-if agency in biological organisms as well as a nexus of other causal factors often conspire to pull us away from memetic explanations, even as informational explanations of other kinds become appropriate.

REFERENCES

Boyd, R., and P. Richerson. 1985. *Culture and the Evolutionary Process*. Chicago: University of Chicago Press.

Cao, R. 2012. "A Teleosemantic Approach to Information in the Brain." *Biology and Philosophy* 27 (1): 49–71.

Dawkins, R. 1976. *The Selfish Gene*. Oxford: Oxford University Press.

———. 1993. "Viruses of the Mind." In *Dennett and His Critics: Demystifying Mind*, ed. Bo Dahlbohm. Oxford: Blackwell.

Dennett, D. C. 1978. *Brainstorms*. Montgomery, VT: Bradford.

———. 2017. *From Bacteria to Bach and Back*. New York: Norton.

Fracchia, J., and R. C. Lewontin. 1999. "Does Culture Evolve?" *History and Theory* 38:52–78.

Godfrey-Smith, P. 2009. *Darwinian Populations and Natural Selection*. Oxford: Oxford University Press.

Henrich, J. 2015. *The Secret of Our Success: How Culture Is Driving Human Evolution*. Princeton, NJ: Princeton University Press.

Kim, J. 1993. "The Myth of Nonreductive Materialism." In *Supervenience and Mind*, ed. J. Kim, 265–84. Cambridge: Cambridge University Press.

Lewontin, R. C. 1985. "Adaptation." In *The Dialectical Biologist*, ed. R. Levins and R. C. Lewontin. Cambridge, MA: Harvard University Press.

Oyama, S., P. E. Griffiths, and R. D. Gray, eds. 2001. *Cycles of Contingency*. Cambridge, MA: MIT Press.

Richerson, P., and R. Boyd. 2005. *Not by Genes Alone: How Culture Transformed Human Evolution*. Chicago: University of Chicago Press.

Sperber, D. 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. R. Aunger, 163–73. Oxford: Oxford University Press.

Sperber, D., and N. Claidière. 2006. "Why Modeling Cultural Evolution Is Still Such a Challenge." *Biological Theory* 1 (1): 20–22.

Sterelny, K. 2006. "Memes Revisited." *British Journal for the Philosophy of Science* 57:145–65.

Ullman-Margalit, E. 1978. "Invisible Hand Explanations." *Synthese* 39 (2): 263–91.

Waters, C. K. 2007. "Causes That Make a Difference." *Journal of Philosophy* 104 (11): 551–79.

Wheeler, J. A. 1989. "Information, Physics, Quantum: The Search for Links." Presented at the 3rd International Symposium on the Foundations of Quantum Mechanics, Tokyo.

Wimsatt, W. 1999. "Genes, Memes, and Cultural Heredity." *Biology and Philosophy* 14:279–310.

https://doi.org/10.1086/710518 Published online by Cambridge University Press