

Is Cultural Fitness Hopelessly Confused?

Grant Ramsey and Andreas De Block

ABSTRACT

Fitness is a central concept in evolutionary theory. Just as it is central to biological evolution, so, it seems, it should be central to cultural evolutionary theory (CET). But importing the biological fitness concept to CET is no straightforward task—there are many features unique to cultural evolution that make this difficult. This has led some theorists to argue that there are fundamental problems with cultural fitness that render it hopelessly confused. In this essay, we defend the coherency of cultural fitness against those who call it into doubt.

- 1 *Introduction*
 - 2 *Cultural Fitness is Not Superfluous*
 - 3 *The Individuation Problem*
 - 3.1 *The generations problem*
 - 3.2 *Cultural growth versus cultural reproduction*
 - 4 *The Repeated Learning Problem*
 - 4.1 *The retention problem*
 - 4.2 *Is fitness transitive?*
 - 5 *Towards a Concept of Cultural Fitness*
 - 5.1 *Biological and cultural individuals*
 - 5.2 *Mutation and transmission biases*
 - 6 *Conclusions*
-

1 Introduction

The 1970s witnessed a grand effort to ‘biologize’ human behaviour. In his magnum opus, *Sociobiology*, E. O. Wilson ([1975]) famously (or infamously) included a twenty-seventh chapter in which he applied the sociobiological framework he had developed for ants and other non-human animals to *Homo sapiens*. This quest to biologize human behaviour, to attempt to

explain it using the same models used for other non-human animals, received intense and immediate censure. Many cultural anthropologists saw such models as not only too crude to capture the richness of human culture, but also empirically wrong (for example, Sahlins [1976]; Barkow [1978]): We just don't act to promote our biological fitness. Instead, there are many human acts that are biological dead ends—the practice of celibacy, for example. These behaviours seem to undermine the reproductive success of individuals, nonetheless, many individuals adopt these cultural practices.

Although some of the critiques of human sociobiology were off target, it seems correct that a merely gene-focused view of human behaviour will not suffice. One response to this apparent failure of an exclusively gene-focused view was to see humans as possessing two inheritance tracks, not merely one: one track is genetic, the other is cultural. This 'dual inheritance' account was developed by Cavalli-Sforza and Feldman ([1981]), Boyd and Richerson ([1985]), Durham ([1991]), and others. Their response constituted not a rejection of biological models, but their extension to the cultural realm. Along with this came questions about the degree to which such an extension requires revisions of the models and modifications of central concepts. Despite the substantial progress in cultural evolutionary theory (CET), there remains much debate and misunderstanding about some of the core concepts at the foundation of this theory.

The focus of this essay will be on one of these core concepts, that of cultural fitness. The driving question is whether the concept (or family of concepts) of fitness, developed within the context of biological evolution, can be applied to cultural evolution and, if so, how it should be understood within this new context. In accounts of cultural evolution, cultural fitness is a concept that seems to do considerable work. Many theorists have stressed that the explanatory power of CETs depends on a well-defined concept of cultural fitness. For example, according to David Sloan Wilson:

The ability to define fitness independently of what evolves saves the concept of natural selection from being a tautology. For the meme concept to escape the same problem, we must define cultural fitness independently of what evolves. If the first four notes of Beethoven's fifth is a powerful meme only because it is common, we have achieved no insight. ([1999], p. 206)

And Sterelny and Griffiths ([1999], p. 334) subscribe to Sober's ([1992]) critique of CETs, namely, 'we have no explanation of the nature of the fitness of ideas, nor do we typically understand why they differ in fitness'. To be clear, we will not argue that CETs have explanatory power. We will only show that if CETs have any explanatory power, this explanatory power depends—at least to some extent—on a sound concept of cultural fitness.

Calls for a definition of 'cultural fitness' have resulted in numerous proposals. Cultural fitness has been defined as the relative success of cultural

parents in transmitting a given cultural behaviour pattern (Avital *et al.* [1998]), the ability of a cultural variant to be copied, stored, and reproduced more frequently than other cultural variants (Blackmore [2010]), the replicant's capacity for replication (Cavalli-Sforza and Feldman [1981]), and the susceptibility to imitation (Jan [2000]), to give just a few examples. Unfortunately, many of these proposed definitions are not very precise, and it is no surprise that this has led to considerable confusion over cultural fitness. Henrich *et al.* ([2008]) list conceptual problems with cultural fitness as the source of one of the five most fundamental and recurrent misunderstandings about cultural evolution.

Part of the problem with developing a concept of cultural fitness is that there is considerable conceptual confusion and debate surrounding biological fitness (Ariew and Lewontin [2004]). There is uncertainty about how fitness should best be modelled and measured, and there are debates concerning whether it is token organismic fitness or trait fitness that is fundamental to evolutionary theory (Sober [2013]). Reviewing the debates concerning precisely how to define or model fitness would take us too far afield (for a review of some of these issues, see Ramsey and Pence [2013]). Instead, what is important to note is that one has to be very cautious when borrowing a concept of fitness from biology to serve as a foundation for cultural fitness. In addition to the challenges concerning borrowing concepts from biology—or developing concepts in parallel inspired from biological theory—there are several problems unique to cultural fitness. It has been argued, for instance, that the problems with cultural fitness are so grave that there is little value to trying to create a general theory of cultural evolution (Gabora [2011]; Strimling *et al.* [2009]). In this article, we will address the question of whether the pursuit of a coherent, useful concept of cultural fitness is indeed hopeless. We will use a two-step strategy to challenge key arguments against cultural fitness. First, we will show that the arguments do not actually undercut the use of cultural fitness in CETs. Second, to the degree that this first strategy fails or is insufficient, we will show that very similar problems arise for biological fitness. The second step intends to show that even if the problems are real and difficult to solve, they are not so serious as to warrant the automatic dismissal of evolutionary approaches to cultural explananda since the same reasons would undercut the use of fitness in the biological realm.

The arguments that our two-step strategy is supposed to tackle are the following: (i) cultural fitness is ill defined because cultural generations are ill defined (Section 3.1); (ii) it is not clear what the bearers of cultural fitness are (Section 3.2); (iii) there is no single scalar value for cultural fitness (Section 4.1); and (iv) cultural fitness values are intransitive (Section 4.2). Although these critiques are not an exhaustive list, we feel that they constitute the most difficult challenges to the tenability of cultural fitness.

Once we have defended cultural fitness from these critiques, we turn to the question of what a positive account of cultural fitness might look like (Section 5). Before we can mount our defence of cultural fitness, however, we first have to explain why one cannot solve the problems we raise by just eliminating cultural fitness from CET.

2 Cultural Fitness is Not Superfluous

Before we offer a defence of cultural fitness, we should first consider whether cultural fitness is an important—or even necessary—component of CET, since no defence is needed if cultural fitness is superfluous. Certainly, not everyone is convinced that cultural fitness is an important component of CET.

Lewens ([2012]), for example, argues that eschewing cultural fitness may be the best way to proceed concerning the problems that have been raised with regard to cultural fitness. This solution is not entirely new. In the 1989 Nancy L. Schwartz Lecture, Reinhart Selten sketched a fictitious dialogue between a population geneticist and an adaptationist about the promises and perils of CET. Both the population geneticist and the adaptationist claim that cultural fitness is not a concept they need for studying cultural evolution. According to Selten's adaptationist, 'cultural evolution tends to the maximization of biological fitness. We do not need a concept of cultural fitness' (Selten [1991], p. 12). Selten's population geneticist disagrees with the reason for rejecting cultural fitness, but he agrees that 'it is unclear whether a cultural fitness concept could be useful' (Selten [1991], p. 12). The population geneticist then goes on to argue that models of cultural evolution are mostly models of quantitative inheritance and that the fitness concept is not useful for quantitative inheritance. Much like Lewens, Selten's population geneticist believes that mathematical and statistical tools suffice to explain population-level phenomena as the aggregated effects of interacting individuals.

Suggestions that we go without fitness have been voiced in biology as well. Dawkins ([1982]), for instance, devoted a whole chapter in *The Extended Phenotype* to argue for the view that 'fitness is a very difficult concept, and that there might be something to be said to do without it whenever we can' (Dawkins [1982], p. 179). According to Dawkins, natural selection is too valuable a theory to let it be compromised by what he clearly sees as philosophical quibbles over an unnecessary concept. Perhaps much to Dawkins's surprise, some philosophers seem to share his view that theories of evolution do not need a fitness concept. Millstein, for instance, has claimed that one can avoid disputes over fitness in evolutionary biology because it is possible to define evolution by natural selection in general terms without invoking fitness: selection is a 'discriminate sampling process whereby physical differences between organisms are causally relevant to differences in reproductive success'

(Millstein [2006], p. 640). Indeed, by substituting ‘causally relevant physical differences’ in place of ‘fitness differences’, she has successfully avoided the use of the term ‘fitness’. Although there is in principle nothing wrong with this definition of selection, it seems to show only that one can avoid the term ‘fitness’, but it does not show that one can avoid the concept in reasoning about selection. After all, the phenotypic differences between organisms that are causally relevant to differences in reproductive success are—or at least give rise to—fitness differences. Moreover, the measurement of the expected influence of an individual on the presence of the trait in the next generation requires a measure of the relevance of the physical differences for reproductive success. And this measure is the good that fitness models can and should deliver. The conceptual link between selection and fitness has not gone unnoticed by biologists. In fact, it is not uncommon to define selection in terms of fitness differences. Lewontin ([1970]) lists fitness differences as one of the necessary requirements for evolution by natural selection, and some standard evolutionary textbooks define selection as occurring via consistent differences in fitness (for example, Futuyma [1998]).

A similar counterargument can be developed against proposals to do away with ‘cultural fitness’. Of course, one can argue that cultural evolution models can do without ‘fitness’ because they try to capture the tendency of a cultural variant to spread through the population due to its effects on the behaviour of its adopters. But this tendency is exactly what some theorists have called ‘cultural fitness’, so this counterargument would only have shown that we could do without the term ‘cultural fitness’, which is not an important theoretical achievement. It does nothing to show that the concept of fitness is superfluous. On the contrary, selection cannot do the theoretical work it is supposed to do without any fitness concept. This is as true for biological evolution as it is for cultural evolution. Both in evolutionary biology and in CET, selection and fitness are conceptually linked; and one can avoid the term fitness, but not the concept (see also Pence and Ramsey [2013]), at least if one holds that selection has an important explanatory value for CETs.

Lewens’s eliminativist attitude towards cultural fitness stems from his view that CETs can do without selection. More specifically, Lewens has argued that a form of population thinking is the core of dual inheritance theory, and that this form of population thinking does not need natural selection: ‘Natural selection although present in some of Boyd and Richerson’s models, is not their explanatory centrepiece; population thinking is’ (Lewens [2010], p. 833). Although it is true that if CET does not need selection, it is not apt to need fitness either, we agree with Houkes ([2012]) that the explanatory value of dual inheritance theory depends largely on its selectionist character: natural and cultural selection of cultural items occupy a central place in many of the explanations that CETs have to offer. A non-selectionist CET would not only do

away with fitness, it would also do away with much of the explanatory power of CETs. This is not to say that the only explanations and predictions of CET derive from cultural fitness. On the contrary, there are important results that derive from population thinking and the application of statistical models to cultural trait dynamics. Nevertheless, we hold that cultural fitness and the process of natural selection due to fitness differences is central to CET. Not only do Boyd, Richerson, and their colleagues argue extensively for the natural selection of social transmission biases (Boyd and Richerson [1985]; Henrich and Boyd [1998]), the biased transmission of cultural variants itself can be understood as cultural selection (Cavalli-Sforza and Feldman [1981]), as we will explain in Section 5. Just as perceptual biases are often considered to be components of selection in sexual selection (Ryan [1990]), we suggest that the transmission biases should be a part of selection/fitness in cultural evolution.

Contra Lewens and others, we think that CET requires a respectable concept of cultural fitness in order to be a viable theory of cultural change. But even if our view is incorrect and a viable CET does not completely depend on a respectable notion of cultural fitness, few would disagree that it would be a good thing to have such a notion. The following sections try to demonstrate that it is possible to develop such a respectable notion because the putative fatal problems of cultural fitness are not in fact fatal.

3 The Individuation Problem

Cultural transmission differs in many respects from genetic transmission. For example, one's cultural parents are not always one's biological parents. One can learn socially from teachers with whom one is biologically only very distantly related. Older people can learn from their children, and a lot of cultural information is transmitted between peers. Genes, however, are almost always transmitted from parents to their biological offspring only.¹ This difference is emphasized in all introductory texts about cultural evolution (for example, Richerson and Boyd [2005]; Mesoudi [2011]). In these introductions, cultural transmission is usually understood as transmission between distinct biological organisms, as if cultural individuals are biological organisms. But is this true? There seem to be cases that cast doubt on this assumption (Jablonka [1994]). For example, it appears possible that one (biological) individual can be its own cultural parent: people can, for example, learn things from their diary that

¹ We are taking animals as our prototypical organism. This, of course, is to ignore prokaryotic life, in which sex (exchange of genes) is not necessarily linked with reproduction. We will merely note here that we will not focus on prokaryotic life and that there are more parallels with prokaryotic genetic dynamics and cultural dynamics than there are with eukaryotic genetic dynamics (Jablonka and Lamb [2005]). But to expand on these parallels and contrasts is well beyond the scope of this article.

they had forgotten about. Furthermore, cultural parents can still transmit information to their cultural offspring even after their death as a biological organism. And if you read co-authored papers like this one, you seem to have one teacher consisting of two biological organisms. These cases illustrate that it might be difficult to get a clear understanding of what, precisely, cultural individuals are, and without a clear conception of cultural individuals, it is difficult to know how one can clearly conceive of cultural fitness. This individuation problem can be broken down into two distinct questions: First, what counts as a cultural generation? Second, what distinguishes growth from reproduction? We will consider these questions in turn.

3.1 The generations problem

Some have argued that cultural fitness is ill defined because generations are not well defined (Gabora [2011]). If fitness requires the existence of generations, then the absence of clear-cut generations in culture entails that the cultural fitness of a variant cannot be determined at all. As Gabora argues:

In biological terms, the fitness of an organism is the number of offspring it has in the next generation. The term 'generation' is applicable when individuals are irretrievably lost from a population and replaced by new ones. But with respect to culture, an outdated (or seemingly 'dead') idea or artifact can come back into use (or seemingly 'come back to life') when styles change or circumstances become right. Because there is no hard and fast distinction between a living entity and a dead one, there is no basis for determining what constitutes a generation. Thus the term generation does not apply to culture. ([2011], p. 74)

Although we think that Gabora's argument is not very conclusive, we readily admit that there is a problem with defining generations in cultural evolution. But even if this problem is intractable, or even if we had to conclude that there is no such thing as a cultural generation, we could still hold on to the view that cultural fitness is not doomed. This is true because fitness—biological or cultural—can be a useful concept with or without generations.

It is true that the biological fitness of an individual organism is sometimes defined as the individual's relative contribution to the gene pool of future generations of a population (Fisher [1999]; Haldane [1924]). Understandably, some cultural evolutionists have used this definition as a template for the definition of cultural fitness, which they then see as the extent to which an individual can affect the proportional representation of a cultural trait in the next generation (Wenseleers *et al.* [2010]). Other cultural evolutionists, in contrast, avoid any reference to generations when they define cultural fitness, perhaps because they are aware of the difficulties with cultural generations. For instance, Henrich ([2004], p. 21) thinks that cultural fitness

‘measures the degree to which a particular value of φ , representing stuff stored in the head of individual i in group j at time t , affects its proportional representation in the population at time $t+1$ ’. Here, time-indexing replaces generation-indexing. This is far from a novelty coming from CET proponents. In biology, fitness is often time-indexed: genetic fitness has been characterized as a ‘measure of the rate at which allele or linkage-group frequencies change within a gene pool through time’ (Wittenberger [1981], p. 614). Time-indexing can be advantageous for biological fitness, and the use of non-generational fitness measures has been defended over and over again (Metz *et al.* [1992]; Coulson *et al.* [2006]).

Time-indexing and generation-indexing, we should emphasize, are not mere differences in model choice that produce no substantive difference. For instance, when generation time substantially varies between individuals of a particular population or when the species under consideration is iteroparous (reproducing multiple times), generation measures become much less useful than time measures (Grafen [1988]). Organisms that reproduce earlier in their life histories will tend to be fitter than others in their population, *ceteris paribus*. Since time-indexing works in such biological cases (and frequently works much better than generation-indexing), then it should work for culture in the absence of (well-defined) cultural generations. We therefore hold that the generation problem for CET is solved by time-indexing.

3.2 Cultural growth versus cultural reproduction

The use of time instead of generations solves one important problem, but there still must be a criterion for counting instances of cultural variants or cultural individuals. What should we count when we measure fitness? Echoing biological fitness definitions, cultural evolutionists give both organism-based (Cavalli-Sforza and Feldman [1981]; Wenseleers *et al.* [2010]) and variant-based (Blackmore [1999]; Henrich [2004]) definitions of cultural fitness. Organism-based cultural fitness is often defined as the probability that the individual organism will be a model for—or be imitated by—others. Cultural variant-based views, on the other hand, focus on the cultural variants (cultural replicators or ‘memes’) and base their fitness on their probability of surviving and reproducing.

To distinguish these two approaches, consider for a moment how we could measure the cultural fitness of this essay. Do we count instances of the paper (number of times it has been downloaded or printed, say)? Or do we count the number of people who have read it, or who have cited it? If a person prints multiple copies, one to have at home and another to keep at the office, does this count as an increase in fitness (variant-based approach) or not (organism-based approach)?

This difficulty touches on a problem that seems particularly pressing for cultural evolution, namely, the problem of distinguishing growth from reproduction. In biology it is clear that counting the number of white/purple flowers of the common foxglove (*Digitalis purpurea*) is not how one should measure the fitness of the white/purple flower variants, whereas it is appropriate to count the plants with white/purple flowers (Janzen [1977]). This is not to say that flower number is irrelevant to fitness—it could even serve as a crude proxy, especially if flower number per plant is relatively constant—only that it is but one component of fitness. What really matters for evolution is the change over time in the frequency of plants with white/purple flowers.

Let's consider how this can help us distinguish cultural growth from cultural reproduction. Suppose that one individual designs a new T-shirt. The idea is in her head, but she also has a few drawings of the new T-shirt in her designer portfolio. Because she was afraid she would lose these drawings, she also scanned them and saved a copy of this scan on her laptop and on a flash drive. Because the T-shirt was thought to become an instant success, it was produced in high numbers. Despite this optimism, only five percent of the production were sold. However, the T-shirt subsequently became popular in the cultural niche of Belgian art house film lovers, all of whom bought twenty T-shirts because a popular Belgian art house director wears them during all his interviews. So what should we count when we want to know the cultural fitness of the T-shirt design? The drawings? All the T-shirts produced or only those that were sold? Or only the individuals who bought a T-shirt?

From a memes' eye view, meme replication is foundational, and cultural fitness is based on the change in meme frequency between t and $t + 1$. Whether or how humans are affected by the meme replication is irrelevant—what matters is that there are more (or fewer) T-shirt memes in existence. From the meme's point of view, humans (or other cultural organisms) are just channels that are used for the meme's replication: 'a scholar is just a library's way of making another library' (Dennett [1995], p. 346). Blackmore ([2010], p. 260) defines the cultural fitness of a melody along similar lines when she asks, 'what is fitness for a short melody? It is the ability to survive and reproduce, which in terms of music means being copied, stored and reproduced more frequently than other melodies'.

We believe, however, that this purely replicator-based account of cultural fitness runs into problems because it renders cultural fitness futile for a scientific theory of cultural evolution and because it doesn't distinguish cultural growth from cultural reproduction. Both problems are intimately intertwined: only a fitness measure that distinguishes cultural reproduction from growth can be relevant for those sciences that deal with cultural evolution. Blackmore's melody example would imply that the fitness of a particular melody would increase by it being copied or stored more than other melodies,

even though it is only one person who makes all the copies of this melody (on media such as CDs, hard drives, memory sticks, or tapes). But counting these copies to measure fitness looks a lot like counting the number of white or purple foxglove flowers to measure the biological fitness of the white variant. Moreover, the serial copying of a particular melody by one individual may have no substantial effects in the social world. Likewise, xeroxing a drawing of a T-shirt one hundred times every day for a year would dramatically increase the fitness of the T-shirt meme, at least according to Blackmore and Dennett, even if it would not alter anything that biologists or social scientists are interested in.

We take these problems with the memetic approach seriously and propose that one instead takes an alternative approach that preserves a cultural growth–reproduction distinction. Our alternative is to count just one variant per biological organism. That is, the currency of cultural fitness is individual organism’s adoption of cultural variants. This does not mean that the number of copies of the T-shirts, or number of images of the shirt posted on Facebook, are irrelevant—more T-shirts or T-shirt images can increase the probability of their adoption—it is that these copies are not what is to be counted in fitness determinations. The focus is on the number of organisms adopting the variants, not the number of variants. Put another way, while meme copy can serve as evidence for cultural fitness, it is not definitive of cultural fitness.

The theoretical advantages of adopting our proposal are several. First, focusing on biological rather than cultural individuals has the benefit of clarity and simplicity: we do not have to split biological individuals up into different cultural individuals, nor do we have to lump different biological individuals together into one cultural individual. While it is true that biological individuality is not free of problems (Janzen [1977]; Bouchard and Huneman [2013]), the problems of attempting to define cultural individuality independently of biological individuality are many times worse. Second, when population geneticists measure gene frequencies in a population, they are not counting the number of copies of genes within each organism (in which case getting fatter would make you fitter). They instead count each organism only once—and the reasons for this are not merely epistemic.

The individual that CETs are referring to when they offer individual-based definitions of cultural fitness tends to be the biological individual *qua* cultural individual (or *qua* cultural model), not some sort of de-biologized cultural individual (for an interesting exception, see Cavalli-Sforza and Feldman [1981]). Most importantly, the choice to count only one variant per biological individual allows for a more subtle and sophisticated CET.² According to

² Note that in arguing that we should not count more than one variant per person, we are not arguing that we cannot count more than one competing variant per person. A ping-pong player may alternately use the shakehand grip and the penhold grip, and thus has adopted both

CETs, culture matters because it has an influence on behaviour (phenotype). Following Galef ([1988]), Richerson and Boyd ([1997]) distinguish communication from culture, and they define communication as ephemeral information that has no long-term impact on the phenotype. Culture, on the other hand, is the kind of information that has lasting impact on the phenotype. It is for this reason that Ramsey ([2013]) takes having a lasting impact on the behaviour of individuals as a necessary condition for information being cultural. Consequently, counting the number of T-shirts (or the photocopies of a drawing) to measure the cultural fitness of the T-shirt (or drawing) would be like counting the sperm cells or eggs of a mammal to measure its biological fitness. Just like ejaculating is not the same as reproducing in biology, the intracranial or intrapersonal proliferation of cultural variants—including artefacts like books, drawings, or tools—is also not the foundation for cultural fitness. Some cultural artefacts—a statue, say—may have a cultural influence over many people over a long span of time. This does not mean that we need to tie cultural fitness to the influence of such artefacts. Instead, cultural fitness can (and should) be measured in terms of individuals adopting variants. Furthermore, as Ramsey ([2013]) notes, cultural artefacts can influence behaviour without any cultural reproduction taking place. We must thus be careful when theorizing the role of artefacts to carefully distinguish genuine cases of cultural transmission from cases of the artefacts merely influencing behaviour.

That said, we do recognize that there may be explanatory projects where it would be useful to track the number of copies of an artefact. For example, a social scientist can be interested in the effect of cigarette use on the likelihood that another person would take up smoking, and on the possibility that each smoked cigarette increases (or decreases) the chance that another person will take up smoking. In cases like these, it could be legitimate to consider measuring fitness by counting cigarettes or counting instances of smoking behaviour. But even then, we think it should be a last resort. After all, there are the epistemic problems of the meme-centred approach we already sketched. Moreover, even in the cigarette example, what seems to matter most is how many individuals adopt the smoking versus the non-smoking variant. Lastly, the number of cigarettes one particular individual smokes can be conceptualized as a particular cultural variant that is adopted by that individual. It should be noted that in biological evolution, the evolution of quantitative traits is dealt with in a similar organism-based way.

variants. Instead, what we are arguing is that we should not count the shakehand grip more than once for one individual, even if the player has photos of them playing with this grip, writes in their diary about playing with this grip, and so on.

4 The Repeated Learning Problem

Even if one has solved the problem of individuating cultural variants, and one can treat cultural variants as properties of organisms in much the same way as genic variants are properties of organisms (capable of being phenotypic trait difference makers, for example), there are features of cultural dynamics unique to culture. One of these features, the impact of repeated learning on cultural evolution, causes deep difficulties with the very possibility of cultural fitness, or so argue Strimling *et al.* ([2009]) (henceforth, SEE). In this section, we analyse SEE's arguments and ultimately conclude that the 'problems' they identify are neither unique to cultural fitness nor fatal to its prospects.

4.1 The retention problem

There are clear ways that genetic and cultural reproductive dynamics differ. Consider four fairly obvious, though important, features of the dynamics of genes in the sexual reproduction of animals: First, one has precisely two genetic parents. Reproductive technology may soon render this false of some humans, but for now this is clearly the case. Second, the genetic constitution of an individual is independent of the properties of the history of parental mating. By this we mean that each mating is basically a genetic rolling of fair dice—the dice are not weighted based on how they have been previously thrown. Third, once an offspring is born, its genetic traits are relatively free from loss or transformation. This is not to deny that mutations can occur, but that mutations are rare and the ratio of mutated genes to non-mutated genes in an individual is approximately zero. Finally, parental mating post-conception has no bearing on one's genetic constitution. If an individual has allele *A*, the later mating behaviour of its parents has no bearing on this fact.

Contrast these four features with those of cultural dynamics: First, with culture, unlike genes, any given trait can have only one or, quite commonly, many parents (Enquist *et al.* [2010]). And the total number of parents for a given organism's set of cultural traits is generally vast: we pick up cultural traits not just from shuffling the traits of two parents (as in the case of our genetic parents), but instead from a multitude of individuals in the population at large (Wimsatt [1999]). Second, because cultural mating can be understood as learning opportunities (genetic mating is coming together to potentially transfer genes in an act of copulation; cultural mating is coming together to potentially transfer cultural variants), while mating history is irrelevant for genes, it matters for culture since the number of opportunities to learn a cultural variant certainly bear on the probability of possessing the variant. Third, once an individual possesses a variant, it could be lost (forgotten) or displaced by a different variant. Fourth, the cultural mating that occurs after a variant is

adopted affects cultural dynamics, since this can lead to the displacement of one variant by another. Finally, the ways that such mating can occur varies considerably depending on the type of culture being passed on. For example, some cultural variants can be passed on immediately after learning (a joke or simple fact, say), but some require years of maturation and/or certification before one can be a cultural transmitter.

These features of cultural dynamics make it clear that there are two central properties of cultural variants that account for at least some of the difference between cultural and genetic dynamics. One is the variant's stickiness (related to what SEE call its 'retention potential'), which is the probability of the variant being retained over time.³ Or, alternatively, one could think of it as the variant's resistance to decay. It is clear that the retention potential for genes is effectively 1. Cultural traits, on the other hand, will have stickiness values between 1 and 0, with intermediate values being the norm.

The other important and unique property of cultural dynamics is the transmissibility of the cultural trait (what SEE call its 'diffusion potential'). This can be understood as the probability that a cultural trait will be adopted given exposure.⁴ Cultural traits that are simple, attractive, or catchy, or that are associated with attractive, powerful, or important models, are apt to have a higher than average transmissibility. For cultural variants, transmissibility can range from 0 (the variant has no probability of being adopted given exposure) to 1 (the variant is guaranteed to be adopted given exposure). With genes, in contrast, the values are generally constant. Any given parental gene has a 0.5 probability of ending up in its offspring, cases of meiotic drive being the rare exception to this rule.⁵ Cultural transmissibility, in contrast, has no default value and is modulated by a diverse array of transmission biases—both direct and indirect (Richerson and Boyd [2005]).

The importance of the fact that cultural variants can take on a full range of stickiness and transmissibility values is that the expected frequency of a variant in a population can dramatically vary depending on the number of learning opportunities: Variants that are sticky but have low transmissibility might

³ There are a number of ways that the 'stickiness' concept could be carved up. For example, one could have a more fine-grained model including variables for the probability of the variant being lost in the absence of competing variants, the probability of being displaced by one type, the probability of being displaced by another type, and so on. These more complex models may well be quite useful, but for our purposes here, a very general notion of stickiness as retention is sufficient.

⁴ As with stickiness, 'transmissibility' is an umbrella term for several related variables. There is the probability of adoption by naïve individuals, the probability of being adopted by an individual with a competing variant of one type, another type, and so on. Again, for the purposes here, these more complex renderings are not necessary.

⁵ Of course, possessing a particular gene can in part be responsible for the individual surviving to reproductive age or being fertile. What is relatively invariant, however, is the conditional probability of the offspring receiving a particular copy of the gene, given that the parent has a copy of it and produces an offspring.

be expected to be common if there are many learning opportunities, but not if there are few. Similarly, variants that have high transmissibility and low stickiness will tend to outcompete less transmissible but stickier variants when there are few learning opportunities. If the variant's fitness is linked to its projected frequency in the population, then fitness is critically dependent on the number of learning opportunities.

What lessons can be drawn from the dependence of a cultural variant's spread on its transmissibility, stickiness, and the number of learning opportunities? One lesson is that if the fitness of a variant is supposed to predict long-term population frequencies, then fitness cannot be equivalent to transmissibility. This is in opposition to some who equate cultural fitness with transmissibility. For example, 'some variants are more likely to be imitated than others, and thus some variants have higher relative "cultural fitness"' (Richerson and Boyd [2005], p. 79), or '[t]here may also be differential fitness, such that some cultural variants are more likely to be transmitted than others and some cultural models (for example, parents and celebrities) are more likely to be copied than others' (Mesoudi [2009], p. 930). These quotes show cultural fitness being tied to probabilities of imitation/transmission and are thus not taking retention into account.⁶ If fitness is supposed to predict the outcome of trait competition, then fitness solely based on copying will perform poorly. Likewise, SEE are correct in concluding:

[...] memetics in its simplest form, where the success of a trait is determined by its ability to 'infect' individuals, works only when individuals have a single learning opportunity. With several learning opportunities, the individual and its previous experiences and choices become important, so individual-based models will be more useful than meme-based models. (Strimling *et al.* [2009], p. 13872)

we agree with SEE that repeated learning and the complications of stickiness and transmissibility cause trouble for the meme-based accounts, but we see such complications merely as a cautionary tale for building a concept of cultural fitness, not as an insurmountable barrier. SEE, as we will now see, hold that these features of cultural fitness lead to insoluble problems.

4.2 Is fitness transitive?

While the complications of stickiness and transmissibility just discussed may not undermine the coherence of cultural fitness, SEE argue that these complications do in fact have a devastating impact on the prospect of cultural fitness;

⁶ One might counter that these quotes are merely taking transmissibility to be one component of fitness. But in support of our interpretation, if one can infer the higher fitness of one variant from its higher probability of being copied, then either there are no other fitness components or the other components are assumed to be equivalent (which is not a realistic assumption).

namely, they can lead to the failure of the transitivity of fitness values. They argue that for cases in which there is a mixed equilibrium of high transmissibility/non-sticky and low transmissibility/sticky types, fitness transitivity can be broken. They produce an example with three variants, 1, 2, and 3, in which $w_1 = w_2$, $w_1 = w_3$, but $w_2 < w_3$. Since if 1 and 2 are equally fit and 1 and 3 are equally fit, one would infer that 2 and 3 are equally fit. But this is decidedly not the case. Consequently, SEE argue that ‘[i]f learning opportunities are few [...] the existence of a success index may be logically impossible, rendering notions of “cultural fitness” meaningless’ (Strimling *et al.* [2009], p. 13870).

First of all, even if this claim is correct, it remains to be seen whether there are in fact many things that we learn socially for which there are only few learning opportunities. Second, and most importantly, although a breakdown of the transitivity of fitness values is *prima facie* troubling, we feel that there are difficulties with their example. One questionable assumption is the inference from fitness equality in one particular ecological set-up (for example, a set-up in which 1 and 2 are the only variants) to the generalization that the variants are equally fit in other ecological set-ups (with another variant, 3, for example, at play). Such inferences are not, in general, warranted (for cultural or for genetic variants). We will suggest that there are cases in the biological realm in which 1 and 2 can exist in a stable equilibrium, 1 and 3 can exist in a stable equilibrium, but that 3 will outcompete 2 or vice versa. Consider the following case.

In the Pacific Northwest of the USA, Douglas fir (*Pseudotsuga menziesii*) is a fast growing tree that dominates many of the forests. While Douglas fir is fast growing and can outcompete other species in sunny patches, it is intolerant of deep shade and fails to grow in the dark depths of a mature forest. Other species, like western hemlock (*Tsuga heterophylla*) or Sitka spruce (*Picea sitchensis*), are slower growing, but are tolerant of shade and can flourish in mature forests. Were there no disturbances, the fir would slowly disappear. But in forests with a moderate degree of disturbance—occasional tree falls, landslides, fires, and so on—the fir can coexist in a stable equilibrium with the shade tolerant hemlock or spruce (Spies and Franklin [1989]). It is thus conceivable that with a moderate degree of disturbance, a forest composed of only fir and hemlock might find a stable equilibrium. Similarly, a forest of only fir and spruce might also stabilize. Using SEE’s framework, one would infer that $w_F = w_H$, $w_F = w_S$, and therefore $w_H = w_S$. We think that it is a bad inference to conclude that the fitness of hemlock and spruce are therefore equivalent. It is quite possible for a forest composed of hemlock and spruce only to come to fixation at one hundred percent spruce or one hundred percent hemlock. Nothing about each species’ competition with the fir says how each will fare against the other. Fitness values simply do not work like this. Transitivity must only hold within a single ecological set-up of species composition and environmental conditions.

If one is concerned that we are making relative fitness comparisons across species instead of among variants of a single species, consider the following example. The failure of transitivity paralleling the tree example above is illustrated by work on the side-blotched lizard (*Uta stansburiana*). This lizard has three distinct morphs, differing in throat colouration and mating behaviour. The mating strategy of the orange-throated males beats the mating strategy of the blue-throated males, but is inferior to the mating strategy of the yellow-throated males. The blue-throated males win when they compete with the yellow-throated males. This means that the fitness of the orange-throated morph is higher than the fitness of the blue-throated morph when these two morphs compete ($w_1 > w_2$). Orange-throated males have a higher fitness than yellow-throated males ($w_2 > w_3$). But it would be wrong to conclude from this that the blue-throated males have a higher fitness than yellow-throated males. Actually, the opposite is true, and there is no reason to think that biological fitness is therefore hopelessly confused; it only shows that biological fitness measures are not transitive across ecological set-ups (Sinervo and Lively [1996]).

We feel that the cases of ecological succession and lizard competition are directly analogous to the case that SEE find unique to cultural fitness. Furthermore, these cases are not particularly troubling for biological fitness and should also not be troubling in analogous cases of cultural evolution. Consider three distinct versions of the same joke. Variant 1 is very difficult to learn, but extremely funny (and therefore unlikely to be forgotten once learned). Variant 2 is moderately easy to learn but only moderately funny. Variant 3 is easy to learn but moderately funny. If each person has on average three opportunities to learn the joke (which is small enough not to have the easily learned jokes dominate, but not so many that the extremely funny version dominates), then it is possible that a stable equilibrium could form with 1 and 2, or with 1 and 3. And since SEE use a stable equilibrium as a criterion for equality of fitness, $w_1 = w_2$ and $w_1 = w_3$. But we would expect the easy joke to dominate the moderately easy one, thus $w_2 < w_3$, and there is nothing strange about this if the competition is only between jokes 2 and 3. The breakdown of transitivity that SEE argue can occur thus does not arise when one is careful not to project fitness values beyond their local ecological set-ups—the transitivity problem is thus a phantom problem.

5 Towards a Concept of Cultural Fitness

This article has focused on defending the tenability of cultural fitness. Although this is the core aim of the article, our defence can aid us in offering something positive about how cultural fitness can be conceived. In this section, we will draw out some of these positive implications for cultural fitness.

5.1 Biological and cultural individuals

As soon as we grant cultural variants their own independent existence, we are faced with myriad difficult questions about their individuation, persistence, and reproduction. As the case of the T-shirt exemplifies, the memeticist has to determine whether backing up a T-shirt design by, say, emailing it to oneself counts as meme reproduction—and if it does, how much reproduction (it is apt to exist on the email server plus all of the email clients that the artist uses). Our solution has been to individuate culture not in terms of independently existing cultural units (or memes), but to piggyback on the individuation of organisms: cultural reproduction occurs just in case an organism has adopted the cultural variant. This view has the advantage that biological organisms are considerably easier to individuate than memes—aphids and dandelions notwithstanding (Janzen [1977]).

One implication of our arguments for an organism-centred account is that such an account is not merely a variant of the meme-based view. This implication is in contrast to some of the claims or assumptions of other authors who see the two views as terminological variants of one another. Henrich *et al.* ([2008], p. 127), for instance, write:

[...] the appropriateness of tracking fitness from the perspective of the meme (assigning fitnesses to alternative memes) or to individuals (or groups) is merely a modelling convenience. For example, just as with genetic evolution, it is not 'more correct' to view fitness in association with memes, individuals, or groups.⁷

Others argue that there is a fact of the matter which description is correct. Brandon and Nijhout ([2006]), making a similar argument in the case of genetic evolution, argue that the genic and genotypic views of fitness are empirically distinct, predicting different dynamics, and that the genic models are empirically wrong.

The organism-centred view can cover much of what we would want of a CET. We are not, however, arguing that a single organism-based fitness concept is all we would need in any context. Some cultural phenomena may call for a group- or meme-level fitness measure. This is analogous to the multi-level fitness concepts developed by Wilson and others (for example, Wilson [1977]; Boyd and Richerson [1990]). And just as a group selection component may play an important role in cultural evolution, there are contexts in which a meme-level approach would be fruitful. In biological evolution, there are cases of meiotic drive in which genes can bias their transmission rates, thereby rigging the meiotic lottery. Just as the gene's eye view is valuable in these cases, so might the meme's eye view be useful in some cases. Thus, our view

⁷ See also (Wimsatt [1999]).

allows for multi-level selection, but we hold that our individual organism-based view is apt to cover much of the ground for CETs, and because of its conceptual and epistemological advantages, it should be the default understanding of cultural evolution.

5.2 Mutation and transmission biases

Given that culture, like genes, is tied to biological organisms, we should now ask how best to extend the full evolutionary framework to culture. Let's begin by considering key components of genetic evolution. First, there is the contribution that the gene makes to the organism's fitness. This can be understood as the organism's fitness with the gene minus its fitness without the gene (or with a different variant of the gene) and is standardly estimated in population genetics by measuring the average effect of a gene in a population (a linear regression of organismic fitness on number of copies of the gene). Second are the chance gene frequency changes due to drift. Third is the mutational component. Fourth is the degree to which the gene biases its chances of making it into the gamete. In genetic evolution, the first can take on a wide range of values: some genes have a strong positive effect, while others lead to sterility or early death. The value of the second is highly dependent on population size. The third generally has a low value, since probabilities of genetic mutations are often low. The fourth is typically zero—the gene does not bias its probability of being selected in the meiotic lottery, though there are cases of meiotic drive in which a gene can cheat the lottery.

How do these four components of genetic evolution align with the components of cultural evolution? We will answer this question by first introducing the framework of Richerson and Boyd ([2005]) and then modifying it to align it with our framework. Richerson and Boyd offer a tripartite division of evolutionary forces: random forces, decision-making forces, and natural selection. Random forces include random changes to cultural variants due to misremembering as well as cultural drift. The set of decision-making forces is a large category that includes guided variation (the active invention or modification of cultural variants) and biased transmission. Biased transmission is, in turn, broken down into three categories: content-biased biases (transmission biased by features of the content like its usefulness or ease of being remembered), frequency-based biases (cultural variants being disproportionately adopted due to being in a high or low frequency), and model-based biases (biases due to features of the model, not the variant or its frequency). Finally, there is natural selection, which is restricted to '[c]hanges in the cultural composition of a population caused by the effects of holding one cultural variant rather than others' (Richerson and Boyd [2005], p. 69).

Let's see how one might map Richerson and Boyd's framework onto the genetic framework. Their random forces map onto the genetic mutations and drift. Their natural selection maps on to the fitness contribution of the genes. Now consider the decision-making forces. The guided variation appears to have no analogue in the genetic framework; there is no evidence that organisms ever synthesize genes to fulfil particular ends (Mesoudi [2008]). Epigenetic transformations come the closest to this, but the analogy is imperfect (Jablonka and Lamb [2005]). It is also difficult to know how to map biased transmission onto the genetic framework. Content-based biases are best understood as analogous to meiotic drive, since the traits are biasing their probability of transmission (Boyd and Richerson [1985]). Model-based biases are closest to genetic hitchhiking (Gabora [1997]), where a gene gets passed on because of the other genes that the organism possesses, though in this case it could be both cultural and genic traits that make the model attractive, thereby increasing its chances of being copied. Finally, frequency-based biases are not unlike frequency-dependent selection, except that Richerson and Boyd place it outside of the domain of natural selection. For Richerson and Boyd, therefore, natural selection is a relatively small player in cultural evolution. Most of the cultural changes we observe—changes in diet, fashion, and so on, and especially those that occur over short time scales—will be due to random or decision-making forces, not natural selection. Cultural fitness, on this view, has a similarly restricted domain if we preserve the tight link between natural selection and fitness.

One of the lessons from our discussion of the stickiness and transmissibility of cultural variants is that linking fitness only with transmissibility (the probability of being copied) makes fitness a poor predictor of the success of cultural variants. Thus, a conception of fitness that is more expansive than (at least some renderings of) memetic fitness is in order. But just as the memetic focus on transmissibility is too restrictive, we also find Richerson and Boyd's view overly restrictive. We hold that there are many ways that a cultural variant can be fit. It can boost the reproductive output of the organisms in which it finds itself (Richerson and Boyd's cultural natural selection), but it can also be highly fit by boosting transmissibility (Richerson and Boyd's transmission biases) as well as stickiness. We thus call for an expansive understanding of cultural fitness, one that will endow it with the full suite of components helping to explain directional evolutionary change. Ideally, this should go together with a more expansive understanding of selection. If we are right, the debate between Richerson and Boyd ([2005]) and Cavalli-Sforza and Feldman ([1981]; see also Durham [1991]) on the notion of cultural selection is more than just a semantic argument. We think that linking fitness to selection is so important that it warrants the use of 'selection' to describe biased transmission (Houkes [2012]). Of course, biased transmission differs from natural selection,

but that is exactly why Cavalli-Sforza and Feldman propose to call these forms of selection ‘cultural selection’.

To summarize our positive contribution to a concept of cultural fitness: First, cultural fitness is measured not by counting all disembodied cultural variants, but by counting only one variant per biological organism. Second, cultural fitness should be able to explain and—to some extent—predict the outcome of cultural dynamics. Most often, the relevant outcome is the proportion of a population of biological individuals that have adopted a cultural variant after a particular time interval. In very specific cases, the relevant outcome can be the relative frequency of a cultural variant or of groups of individuals that have adopted the variant. Third, cultural fitness should track many of the mechanisms and processes that make cultural evolution directional in a cultural lineage. Because natural selection on organisms plays a relatively weak role in increasing the chances of a cultural variant to spread in a biological population, cultural fitness differences should also be based on various biases, such as transmission biases and what one may call ‘retention biases’. Fourth, while an expansive concept of cultural fitness is desirable, one should be careful not to make it too expansive. For example, guided variation is, in general, best not considered to be a component of the fitness of cultural variants. The reason for this is that cultural fitness is used not just for understanding changes in the frequencies of cultural variants, but also for understanding the adaptive features of the cultural variants themselves. A cultural trait can be adaptive because it is ‘intelligently designed’, by which we mean that individual organisms produce the variant with the goal of producing its adaptive nature. Alternatively, it can be adaptive due to a process of blind variation and natural selection—the adaptive features due in part, that is, to their cultural fitness.

6 Conclusions

Cultural fitness is not as hopelessly confused as some authors have contended. If cultural fitness is tied to organisms (and not free-floating memes), then many of the conceptual problems with cultural fitness dissolve. There are still some challenges distinguishing growth from reproduction, but none worse than occur in biology outside of the cultural realm. And while repeated learning challenges the memetic view, it does not ultimately challenge the tenability of cultural fitness. Cultural fitness can do similar work in the study of cultural evolution as biological fitness does in the study of biological evolution. The differences between biological fitness and cultural fitness are real and substantial, but none of these differences diminish the tenability and usefulness of cultural fitness. In fact, the opposite is true: cultural fitness is a

very useful notion precisely because cultural fitness differs from biological fitness.

Funding

Work on this article was supported by and F+ residential scholarship at KU Leuven awarded to Grant Ramsey.

Acknowledgements

We would like to thank Stijn Conix, Stefano Ghirlanda, Stefan Linquist, Charles Pence, Pete Richerson, Pontus Strimling, the members of the Duke Philosophy of Biology Reading Group, and the reviewers for this journal for providing us with valuable feedback on earlier drafts of this article.

Grant Ramsey
Institute of Philosophy,
KU Leuven
Leuven, Belgium
grant@theramseylab.org

Andreas De Block
Institute of Philosophy,
KU Leuven
Leuven, Belgium
Andreas.DeBlock@hiw.kuleuven.be

References

- Ariew, A. and Lewontin, R. C. [2004]: 'The Confusions of Fitness', *British Journal for the Philosophy of Science*, **55**, pp. 347–63.
- Avital, E., Jablonka, E. and Lachmann, M. [1998]: 'Adopting Adoption', *Animal Behavior*, **55**, pp. 1451–9.
- Barkow, J. H. [1978]: 'Culture and Sociobiology', *American Anthropologist*, **80**, pp. 5–20.
- Blackmore, S. [1999]: *The Meme Machine*, Oxford: Oxford University Press.
- Blackmore, M. [2010]: 'Memetics Does Provide a Useful Way of Understanding Cultural Evolution', in F. Ayala and R. Arp (eds), *Contemporary Debates in Philosophy of Biology*, Chichester: Wiley-Blackwell, pp. 255–72.
- Bouchard, F. and Huneman, P. [2013]: *From Groups to Individuals: Evolution and Emerging Individuality*, Cambridge, MA: MIT Press.
- Boyd, R. and Richerson, P. J. [1985]: *Culture and the Evolutionary Process*, Chicago, IL: University of Chicago Press.
- Boyd, R. and Richerson, P. J. [1990]: 'Group Selection among Alternative Evolutionarily Stable Strategies', *Journal of Theoretical Biology*, **145**, pp. 331–42.

- Brandon, R. N. and Nijhout, H. F. [2006]: 'The Empirical Nonequivalence of Genic and Genotypic Models of Selection: A (Decisive) Refutation of Genic Selectionism and Pluralistic Genic Selectionism', *Philosophy of Science*, **73**, pp. 277–97.
- Cavalli-Sforza, L. L. and Feldman, M. W. [1981]: *Cultural Transmission and Evolution: A Quantitative Approach*, Princeton: Princeton University Press.
- Coulson, T., Benton, T. G., Lundberg, P., Dall, S. R. X., Kendall, B. E. and Gaillard, J. M. [2006]: 'Estimating Individual Contributions to Population Growth: Evolutionary Fitness in Ecological Time', *Proceedings of the Royal Society B: Biological Sciences*, **273**, pp. 547–55.
- Dawkins, R. [1982]: *The Extended Phenotype*, Oxford: Oxford University Press.
- Dennett, D. [1995]: *Darwin's Dangerous Idea: Evolution and the Meanings of Life*, London: Allen Lane Press.
- Durham, W. [1991]: *Coevolution: Genes, Culture, and Human Diversity*, Stanford, CA: Stanford University Press.
- Enquist, M., Strimling, P., Eriksson, K., Laland, K. and Sjostrand, J. [2010]: 'One Cultural Parent Makes No Culture', *Animal Behaviour*, **79**, pp. 1353–62.
- Fisher, R. A. [1999]: *The Genetical Theory of Natural Selection: A Complete Variorum Edition*, Oxford: Oxford University Press.
- Futuyma, D. J. [1998]: *Evolutionary Biology*, Sunderland: Sinauer Associates.
- Gabora, L. [1997]: 'The Origin and Evolution of Culture and Creativity', *Journal of Memetics: Evolutionary Models of Information Transmission*, **1**, pp. 1–28.
- Gabora, L. [2011]: 'Five Clarifications about Cultural Evolution', *Journal of Cognition and Culture*, **11**, pp. 61–83.
- Galef, B. G. [1988]: 'Imitation in Animals: History, Definition and Interpretation of Data from the Psychological Laboratory', in T. Zentall and B. G. Galef (eds), *Social Learning: Psychological and Biological Perspectives*, Hillsdale: Lawrence Erlbaum, pp. 29–48.
- Grafen, A. [1988]: 'On the Uses of Data on Lifetime Reproductive Success', in T. H. Clutton-Brock (ed.), *Reproductive Success: Studies of Individual Variation in Contrasting in Contrasting Breeding Systems*, Chicago, IL: University of Chicago Press, pp. 454–71.
- Haldane, J. B. S. [1924]: 'A Mathematical Theory of Natural and Artificial Selection', *Biological Reviews*, **1**, pp. 158–63.
- Henrich, J. [2004]: 'Cultural Group Selection, Coevolutionary Processes and Large-Scale Cooperation', *Journal of Economic Behavior and Organization*, **53**, pp. 3–35.
- Henrich, J. and Boyd, R. [1998]: 'The Evolution of Conformist Transmission and the Emergence of Between-Group Differences', *Evolution and Human Behavior*, **19**, pp. 215–41.
- Henrich, J., Boyd, R. and Richerson, P. J. [2008]: 'Five Misunderstandings about Cultural Evolution', *Human Nature*, **19**, pp. 119–37.
- Houkes, W. [2012]: 'Population Thinking and Natural Inheritance in Dual Inheritance Theory', *Biology and Philosophy*, **27**, pp. 401–17.
- Jablonka, E. [1994]: 'Inheritance Systems and the Evolution of New Levels of Individuality', *Journal of Theoretical Biology*, **170**, pp. 301–9.

- Jablonka, E. and Lamb, M. [2005]: *Evolution in Four Dimensions: Genetic, Epigenetic, Behavioral, and Symbolic Variation in the History of Life*, Cambridge, MA: MIT Press.
- Jan, S. [2000]: 'Replicating Sonorities: Towards a Memetics of Music', *Journal of Memetics: Evolutionary Models of Information Transmission*, **4**, <http://cfpm.org/jom-emit/2000/vol4/jan_s.html>.
- Janzen, D. H. [1977]: 'What Are Dandelions and Aphids?', *The American Naturalist*, **111**, pp. 586–9.
- Lewens, T. [2010]: 'Natural Selection Then and Now', *Biological Reviews*, **85**, pp. 829–35.
- Lewens, T. [2012]: 'Cultural Evolution: Integration and Scepticism', in H. Kincaid (ed.), *The Oxford Handbook of Philosophy of Social Science*, Oxford: Oxford University Press, pp. 458–80.
- Lewontin, R. [1970]: 'The Units of Selection', *Annual Review of Ecology and Systematics*, **1**, pp. 1–18.
- Mesoudi, A. [2008]: 'Foresight in Cultural Evolution', *Biology and Philosophy*, **23**, pp. 243–55.
- Mesoudi, A. [2009]: 'How Cultural Evolutionary Theory Can Inform Social Psychology and *vice versa*', *Psychological Review*, **116**, p. 929.
- Mesoudi, A. [2011]: *Cultural Evolution: How Darwinian Theory Can Explain Human Diversity and Synthesize the Social Sciences*, Chicago, IL: University of Chicago Press.
- Metz, J. A. J., Nisbet, R. M. and Geritz, S. A. H. [1992]: 'How Should We Define "Fitness" for General Ecological Scenarios?', *Trends in Ecology and Evolution*, **7**, pp. 198–202.
- Millstein, R. L. [2006]: 'Natural Selection as a Population-Level Causal Process', *British Journal for the Philosophy of Science*, **57**, pp. 627–53.
- Pence, C. H. and Ramsey, G. [2013]: 'A New Foundation for the Propensity Interpretation of Fitness', *British Journal for the Philosophy of Science*, **64**, pp. 851–81.
- Ramsey, G. [2013]: 'Culture in Humans and Other Animals', *Biology and Philosophy*, **27**, pp. 457–79.
- Ramsey, G. and Pence, C. H. [2013]: 'Fitness: Philosophical Problems', *eLS*, doi: 10.1002/9780470015902.a0003443.pub2
- Richerson, P. and Boyd, R. [1997]: 'Types of Transmission: A Taxonomy of Cultural Inheritance Systems', in P. Weingart, S. D. Mitchell, P. J. Richerson and S. Maasen (eds), *Human by Nature: Between Biology and the Social Sciences*, Mahwah, NJ: Erlbaum, pp. 313–24.
- Richerson, P. and Boyd, R. [2005]: *Not by Genes Alone: How Culture Transformed Human Evolution*, Chicago, IL: University of Chicago Press.
- Ryan, M. J. [1990]: 'Sexual Selection, Sensory Systems, and Sensory Exploitation', *Oxford Surveys in Evolutionary Biology*, **7**, pp. 157–95.
- Sahlins, M. [1976]: *The Use and Abuse of Biology: An Anthropological Critique of Sociobiology*, Ann Arbor, MI: University of Michigan Press.
- Selten, R. [1991]: 'Evolution, Learning, and Economic Behavior', *Games and Economic Behavior*, **3**, pp. 3–24.

- Sinervo, B. and Lively, C. M. [1996]: 'The Rock-Paper-Scissors Game and the Evolution of Alternative Male Strategies', *Nature*, **380**, pp. 240–3.
- Sober, E. [1992]: 'Models of Cultural Evolution', in P. Griffiths (ed.), *Trees of Life: Essays in the Philosophy of Biology*, Amsterdam: Kluwer, pp. 477–92.
- Sober, E. [2013]: 'Trait Fitness Is Not a Propensity, but Fitness Variation Is', *Studies in History and Philosophy of Biological and Biomedical Sciences*, **44**, pp. 336–41.
- Spies, T. A. and Franklin, J. F. [1989]: 'Gap Characteristics and Vegetation Response in Coniferous Forests of the Pacific Northwest', *Ecology*, **70**, pp. 543–5.
- Sterelny, K. and Griffiths, P. E. [1999]: *Sex and Death: An Introduction to the Philosophy of Biology*, Chicago, IL: University of Chicago Press.
- Strimling, P., Enquist, M. and Eriksson, K. [2009]: 'Repeated Learning Makes Cultural Evolution Unique', *Proceedings of the National Academy of Sciences*, **106**, pp. 13870–4.
- Wenseleers, T., Gardner, A. and Foster, K. R. [2010]: 'Social Evolution Theory: A Review of Methods and Approaches', in T. Szekely, A. J. Moore and J. Komdeur (eds), *Social Behaviour: Genes, Ecology, and Evolution*, Cambridge: Cambridge University Press, pp. 132–58.
- Wilson, D. S. [1977]: 'Structured Demes and the Evolution of Group-Advantageous Traits', *American Naturalist*, **111**, pp. 157–85.
- Wilson, D. S. [1999]: 'Flying over Uncharted Territory: A Review of *The Meme Machine* by S. Blackmore', *Science*, **285**, p. 206.
- Wilson, E. O. [1975]: *Sociobiology: The New Synthesis*, Cambridge, MA: Harvard University Press.
- Wimsatt, W. [1999]: 'Genes, Memes, and Cultural Heredity', *Biology and Philosophy*, **14**, pp. 279–310.
- Wittenberger, J. F. [1981]: *Animal Social Behavior*, Boston, MA: Duxbury Press.