

The swashbuckling anthropologist: Henrich on The Secret of Our Success

Ellen Clarke and Cecilia Heyes

All Souls College, University of Oxford, Oxford OX1 4AL, UK

Abstract

In *The Secret of Our Success*, Joseph Henrich claims that human beings are unique – different from all other animals – because we engage in cumulative cultural evolution. It is the technological and social products of cumulative cultural evolution, not the intrinsic rationality or ‘smartness’ of individual humans, that enable us to live in a huge range of different habitats, and to dominate most of the creatures who share those habitats with us. We are sympathetic to this general view, the latest expression of the ‘California school’s’ view of cultural evolution, and impressed by the lively and interesting way that Henrich handles evidence from anthropology, economics, and many fields of biology. However, because we think it is time for cultural evolutionists to get down to details, this essay review raises questions about Henrich’s analysis of both the cognitive processes and the selection processes that contribute to cumulative cultural evolution. In the former case, we argue that cultural evolutionists need to make more extensive use of cognitive science, and to consider the evidence that mechanisms of cultural learning are products as well as processes of cultural evolution. In the latter case, we ask whether the California school is really serious about selection, or whether it is offering a merely ‘kinetic’ view of cultural evolution, and, assuming the former, outline four potential models of cultural selection that it would be helpful to distinguish more clearly.

Keywords Cultural evolution. Cultural learning. Multi-level selection. Cognitive science.

Homo sapiens: Just another species of animal? While we have much in common with our primate cousins, and with the other inhabitants of the living realm, it is also true that people are pretty special. We do all manner of things that other creatures don’t – for example, some of us wear clothes, vote in elections, give hi-fives, or write book reviews. It is also true that we, as a species, have been very prosperous – we inhabit a huge range of different habitats, our population has exploded and we are able to dominate and control most of the creatures who attempt to share our habitats with us. In *The Secret of our Success* (henceforth SOOS) Joseph Henrich argues that these properties are connected – that those idiosyncrasies of our behaviour that we tend to term our ‘culture’ are responsible for the success our species enjoys. We think this a very reasonable explanation for why, despite humankind’s physical weakness, slowness and dullness of senses, we have been able to out-compete so many much stronger, faster and physically less vulnerable animals. A solitary, unequipped and uneducated human is generally a fairly unimpressive specimen. But allow her access to her culture and all that it can do and she might move mountains.

The devil is in the detail of this appealing thesis, however. What is this ‘culture’, how did we come by it, and what is the nature of its connection to our success? SOOS presents as a trade book, so it is packed with colourful examples, personal observations, anecdotes, and connections with current events; the language is racy and informal. This makes it a lively, engaging read in which the author emerges – in a good way – as a swashbuckling anthropologist. Navigating the high seas of fieldwork, and thrusting the sword of laboratory intervention, Henrich is trying to bring it all together into a new, coherent picture of human origins. But, perhaps inevitably, the details and the logical structure are somewhat difficult to divine. SOOS succeeds in giving an exciting glimpse of the rapidly expanding academic literature on cultural evolution and will be sure to attract new enthusiasts to the topic. A drier text may have served better as a contribution to the literature, however. We accordingly focus, in this review, on discussing how we would *like* to see the field develop, and what we *wish* the book had spelt out. First, we offer the following summary of the main ideas on display.

Henrich claims that human beings are unique – different from all other animals – in that they engage in a process of Cumulative Cultural Evolution (hereafter CCE). It is the technological and social *products* of CCE, he argues, that explain our extraordinary success. Henrich allows that *some* non-human animals may be capable of *some* cultural evolution, but he insists that only humans have become able to sustain cultural evolution in its cumulative form.

A cultural evolution process is *cumulative* when improvements are multi-generational, in that a product used by one person is improved by someone from a subsequent generation, and then improved again in another, later generation. Then, the products are better than anything that can be achieved by a single person in her own lifetime, using trial and error or creative insight.

Crucially, what makes Henrich’s account of our success surprising or ‘secret’ is that it challenges the orthodox Enlightenment assumption that our intrinsic rationality is what separates us from the beasts. Henrich claims we’re not all that smart after all, and the extent to which our genes have been selected to tailor our minds to our inanimate environments is much less than often assumed. It is our *culture* that gets the credit for our apparent brilliance. But the process that gave us our culture is blind – culture emerges as an unintended consequence of interactions between socially learning, but not especially smart, minds. Culture is smarter than we are, and its products ‘work’ without us needing to know how or why they work, however much we might like to think otherwise.

Henrich inherited from the ‘California School’ of cultural evolutionists¹ a commitment to CCE having generally occurred in a group context. And the dynamics of CCE are intended to be populational, so that technological success is sensitive to group size and connectedness, as well as to the rate of change of the environment. One consequence of this is that the accumulated value of culture is rather fragile, being easily lost during environmental shocks or disturbances to group stability.

Another commitment Henrich inherited is to dual inheritance theory. This states that there has been feedback between cultural evolution and genetic selection, wherein the products of CCE have altered *Homo*’s environment such that there has been a genetic response. In other words, there has been standard genetic selection for genes that are well-suited to the culturally altered environment. Henrich presses this point home with many examples of genetically underpinned human anatomical and physiological adaptations whose selection can be understood as occurring within a culturally

¹ Headed by Robert Boyd and Peter J Richerson, as opposed to Dan Sperber’s ‘Paris School’.

altered environment. For example, Henrich pinpoints a mutation which disrupted the development in *Homo* of the large jaw muscles our ancestors presumably used to grind raw food (Chapter 5). The suggestion is that upon the advent of fire-use and other culturally transmitted food processing techniques this gene no longer provided humans with an advantage, and culturally-driven genetic selection caused the disruptive gene to spread throughout the species.

One of the strengths of SOOS is Henrich's emphasis on non-genetic innovation as a possible driver of adaptive evolution, although we wouldn't follow him in viewing this as unique to humans. No view which assumes a unidirectional causal arrow from genes to the environment stands up against the many examples nature provides of reciprocal feedback loops between selection and the environment (Laland et al 2013). Henrich cites the well-known case of the evolution of lactose tolerance, but we can add the earthworm's construction of its underground habitat, and the runaway sexual selection which produced the peacock's tail. The fact that human behaviour has had downstream consequences for the selection of human genes does not, therefore, make us unique, although Henrich makes the further, interesting bet that culture now constitutes the *dominant* selection pressure for humans (p. 57).

In what follows we present two sets of concerns. The first relates primarily to Henrich's account of the *cognitive* processes involved in cultural evolution, and the second set to his picture of the *evolutionary* processes. In both cases we believe that there are more options than Henrich acknowledges, that it is often unclear which of several options he does endorse, and that there is not enough justification offered for those endorsements he does make.

Cognition

Cognition is central to Henrich's analysis. Some of his core claims are about how the human mind works, how it came to work that way, and how this kind of working has resulted in us living the sorts of lives that we lead. The secret of our success lies in our "collective brains". Humans have a set of genetically evolved psychological adaptations that enable us to learn from others in ways that other animals cannot. We have special kinds of social learning called "cultural learning", while they have only the garden varieties (p. 297). Ultimately, "It's our collective brains operating over generations, and not the innate inventive power or creative abilities of individual brains, that explain our species' fancy technologies and massive ecological success" (p. 212). So, cognition – how the mind works - is central to Henrich's analysis, and yet cognitive science is almost completely absent. A large proportion of the book, perhaps as much as a third, describes psychological experiments. We are given lively descriptions of many behavioural studies of chimpanzees, children and people from a range of cultures, and from time to time these are backed up with a bit of neuroscience. But the interpretive framework is always folk psychological. It never goes beyond what the whole agent thinks, wants or decides at a given moment to the kind of sub-personal, information processing analysis that is characteristic of cognitive science. Henrich ignores without comment the kind of research that has dominated psychology since the 1970s, research investigating the computations performed by parts of the mind, and how those parts work together to produce behaviour (Dennett, 1978; Frankish & Ramsay, 2012; Shallice & Cooper, 2011). Cognition is central, but cognitive science is absent without leave.

What is cultural learning?

One consequence of Henrich's reluctance to enquire within - to consider sub-personal psychological processes - is that although we are told that cultural learning is what makes humans special, we never find out exactly what cultural learning is, why it is thought to be distinctively human, or how it makes us special. The most direct attempt to characterise cultural learning is in an early passage:

"Throughout this book, *social learning* refers to any time an individual's learning is influenced by others, and it includes many different kinds of psychological processes. *Individual learning* refers to situations in which individuals learn by observing or interacting directly with their environment, and can range from calculating the best time to hunt by observing when certain prey emerge, to engaging in trial-and-error learning with different digging tools. So, individual learning too captures many different psychological processes. Thus, the least sophisticated forms of social learning occur simply as a by-product of being around others, and engaging in individual learning. For example, if I hang around you, and you use rocks to crack open nuts, then I'm more likely to figure out on my own that rocks can be used to crack open nuts because I'll tend to be around rocks and nuts more frequently and can thus more easily make the relevant connections myself. *Cultural learning* refers to a sophisticated subclass of social learning abilities in which individuals seek to acquire information from others, often by making inferences about their preferences, goals, beliefs or strategies and/or by copying their actions. When discussing humans, I'll generally refer to *cultural learning*, but with non-humans and our ancient ancestors, I'll call it *social learning*, since we often aren't sure if their social learning includes any actual cultural learning." (p. 12-13).

Taken at face value, this passage draws several lines between cultural learning and (other) social learning: cultural learning is "sophisticated", it is not merely a "by-product of being around others", and it is active – "individuals seek to acquire information from others". "Sophisticated" implies complex, but since we are not given any characterisation of the processes underlying individual, social or cultural learning, it is impossible to guess on what dimension(s) cultural learning processes might be especially complex. The by-product and active seeking lines are more promising conceptually but they are in danger of putting many dumb animals in the culture club. For example, before consuming their first meal of solid food, weanling rats seek out a feeding adult member of their colony, and eat whatever that adult is eating (Galef 1971). This enables the weanlings to learn what is safe and nutritious to eat in their local environment. Thus, the weanlings engage in active social learning that depends crucially on what the adult rat is doing, it is not merely a by-product of proximity to adults, but we doubt that Henrich would want the weanlings' behaviour to be an example of cultural learning. By his lights, a psychological capacity that humans share with a wide range of other animals is unlikely to be the secret of our success.

Perhaps, rather than attempting to demarcate cultural learning from other kinds of learning, the quoted passage is simply saying that most cultural learning involves acquiring information from others by mindreading ("making inferences about their preferences, goals, beliefs or strategies") and by body movement imitation ("copying their actions"). This is the view of many experts on cultural evolution (e.g. Fogarty, Rendell & Laland 2012; Tomasello 1999; 2014) but it is problematic (Heyes, 2016a). Unlike active social learning, mindreading and imitation are good candidates to be the secret of our success in that they appear to be distinctively human. But it is far from clear that they are good ingredients for cultural learning.

Why would social learning via mindreading and imitation be more likely than other sorts of social learning to make culture cumulative? This question has not been addressed by any cultural evolutionists, including Henrich, and it is genuinely puzzling. At first blush, mindreading looks like just the wrong kind of process for ‘blind’ cultural inheritance. A learner who is divining a model’s beliefs and desires – working out what an expert wants to achieve, and what she knows about the instrumental power of her actions – doesn’t appear to be adopting behaviour without asking what it is for or why it works. Similarly, body movement imitation may be essential for the cultural inheritance of certain ritual practices and communicative gestures, but the inheritance of only a small fraction of cultural traits depends on encoding the topography of observed body movements. Speech production and comprehension require the learner to represent sounds, not the movements producing those sounds. In the case of tool manufacture and tool use, what matters is how objects interact over time – at what angle, and with how much force, the hammer stone impacts the tool stone to produce a flake of a particular size and shape. As long as the learner can derive this kind of object-related information from watching an expert, they can work out for themselves, later through practice, which body movements will do the job (Heyes 2013).

Whether the quoted passage is an effort to demarcate or simply to define cultural learning by ostension, one of its troublesome consequences is that it leaves many of Henrich’s leading examples of human cultural learning out in the cold. For example, Chapter 4 surveys evidence that human social learning is selective. When we have a choice of models, we are apt to learn from older rather than younger individuals (age bias), from the majority rather than the minority (conformist bias), from successful rather than unsuccessful people (success or payoff bias), and from individuals who are more similar to ourselves (self-similarity bias). As the title of the chapter indicates, “How to make a cultural species”, these biases are meant to be prime examples of cultural learning at work, but each bias has been observed in a range of nonhuman animals (Laland 1996; Hoppitt & Laland 2013), and there is no reason to believe that any of these biases typically involve especially active learning processes, mindreading, or imitation. Indeed, the experimental evidence from animals, children and adults suggests that, in the vast majority of cases, these biases are due to bog-standard psychological processes; to the same attentional mechanisms and associative processes that enable isolated animals to learn about their inanimate environments (Heyes in press; Heyes & Pearce 2014).

In other words, some models are more likely than others to capture our attention because they are bigger (older, majorities), closer to goodies (more successful), or more familiar (self-similarity) than the alternative models, and we are more likely to learn – via standard, domain-general mechanisms – from models to which we attend than from models we ignore. There may well be exceptions to this picture, cases in which humans don’t merely learn from the models who capture their attention, but absorb information from two or more models and then engage an active process, perhaps even mindreading, to decide which of the models to copy. But to find these exceptional cases – what one of us has called “metacognitive social learning strategies” (Heyes 2016b) – you need to enquire within, to ask what is going on in learners’ minds. Reluctant to do this, Henrich simply assumes that all human social learning is cultural learning, and, in the process, violates his own definition of cultural learning.

Where does cultural learning come from?

Lacking a more specific characterisation of cultural learning, let's use the term as short hand for "all features of human psychology that contribute to making culture cumulative". Henrich believes there are many such features, from the social learning biases discussed above, to "pedagogical adaptations" that make us receptive to teaching, via mindreading, overimitation, and norm psychology (see Table 5.1 in Chapter 5). And he insists that all of these features are genetic adaptations. He allows that they may have got started through learning, and that they may still be "tuned" or "calibrated" by experience in early development (pp 52-53), but he is emphatic that they have all been genetically assimilated:

"natural selection, acting on genes, has shaped our psychology in a manner that generates nongenetic evolutionary processes capable of producing complex cultural adaptations. Culture, and cultural evolution, are then a consequence of genetically evolved psychological adaptations for learning from other people. That is, natural selection favoured genes for building brains with abilities to learn from others" (p. 34).

Henrich's insistence that cultural learning is a set of psychological instincts is puzzling for two reasons. First, it makes Henrich's position very similar to that of the High Church Evolutionary Psychologists he claims to oppose. It's true that the High Churchers – including Steven Pinker, Leda Cosmides and John Tooby - are more preoccupied than Henrich with psychological instincts that enable individuals to get information for themselves, rather than from other people. However, social contract reasoning and language, the flagships of Evolutionary Psychology, are social faculties, and the whole idea of "genetically evolved psychological adaptations" comes straight from the High Church hymn sheet.

Yet more puzzling, Henrich's headline commitment to psychological instincts is at odds, not only with many of his passing comments (e.g. "we culturally learn from whom to learn" (p.43); "Acquiring languages alters aspects of our psychology and endows individuals with new cognitive abilities" (p.259)), but with the central theme of his book - the power of cultural evolution to produce complex adaptations. *If, as Henrich claims, cultural evolution can produce complex artefacts, practices and belief systems, why doesn't he consider the possibility that cultural evolution also produces cognitive processes?* On this view, the features of social learning that promote cumulative culture have been made fit for that purpose by cultural evolution, and each child uploads these features through social interaction in the course of development. Thus, our psychological adaptations for cultural inheritance are cognitive gadgets rather than cognitive instincts (Heyes, 2012; in prep).

There are hints in SOOS of three potential responses to the question in italics above:

1) The idea that psychological adaptations are products of cultural evolution implies commitment to a defunct innate-acquired distinction (p. 339, note 34). There are certainly untenable versions of the innate-acquired distinction. However, it is unclear why anyone might think that an untenable version is invoked by the idea that psychological processes are shaped by cultural rather than genetic evolution, but *not* by Henrich's proposal that artefacts, behaviour and beliefs are shaped by cultural rather than genetic evolution.

2) All adaptations become genetically assimilated. If this were true, it would surely justify dismissal of the possibility that psychological adaptations for culture are cognitive gadgets rather than

cognitive instincts, but there is no reason to assume that all cultural adaptations become genetically assimilated. Indeed, Henrich acknowledges, in a typically evocative phrase, that “cultural evolution can sap the strength of [genetic] selection” (p. 92). Cultural adaptations can do a job so well that they prevent a genetic response. For example, the invention of cheese and yoghurt, lactose-light preparations that carry many of the nutritional benefits of whole milk, put a brake on gene-based selection for lactose tolerance.

3) The evidence favours cognitive instincts over cognitive gadgets. This response is suggested by Henrich’s recurrent tendency to invoke as evidence of genetic assimilation any sign that a human psychological attribute a) develops before the age of five years, b) is less impressive in chimpanzees than in children, or c) has an identifiable neural substrate. But these lines of evidence are far from compelling. a) Early development, even in infancy, can indicate that an attribute is being canalised by the cultural environment – by other people – rather than by the genes. b) Children may be better social learners than chimpanzees, not because they have genetic adaptations for cultural learning, but because they have more powerful domain-general mechanisms of learning and culture-soaked developmental environments. c) All psychological processes, whatever their provenance, are implemented in the brain.

The cognitive instincts versus cognitive gadgets debate is about the size and shape of the genetic ‘starter kit’ for cultural learning. Following High Church Evolutionary Psychologists, Henrich and other cultural evolutionists assume that we genetically inherit what might be called ‘whole-big-special’ psychological attributes: genetic resources that amount to programmes or blueprints (whole) for the development of powerful (big), distinctively human (special) cognitive mechanisms, such as selective social learning, mindreading, imitation, and pedagogy. In contrast, the cognitive gadgets view suggests that, although most humans end up having one or another variant of these mechanisms, the genetically inherited resources that contribute to their development are partial, small, and ordinary. Even the most influential genetically inherited resources nudge rather than programme human cognitive development (partial), closely resemble the resources inherited by our recent ancestors, including chimpanzees (small), and differ from those ancestral endowments in ways that would typically be described as quantitative rather than qualitative (ordinary). On the cultural gadgets view, then, the genetic starter kit for cultural learning gives humans a socially tolerant temperament, attentional biases favouring other agents, and uniquely powerful domain-general mechanisms of learning, memory and cognitive control. Guided by cultural evolution, and using these ingredients, social interaction in the course of development constructs dedicated, distinctively human mechanisms for cultural learning. We have domain-specific psychological adaptations for culture but they are culturally rather than genetically inherited (Heyes, 2012; in prep).

Whole-big-special components of cultural learning certainly *could* have been genetically assimilated. As Henrich points out, there are signs that many physiological and morphological adaptations for culture have undergone genetic assimilation, and, if he is right that culture goes back two million years, there has surely been enough time. So, there is nothing inherently far-fetched or implausible about the idea of cognitive instincts for culture. It’s just that the evidence does not point in that direction. Take imitation, one of the first capacities put forward as a crucial component of cultural learning. There would be reason to believe that imitation has been genetically assimilated if there was evidence in humans alive today of “poverty of the stimulus” (Chomsky, 1975), signs that the

development of imitative ability outstrips opportunities to learn to imitate, and/or evidence that variation in imitative ability is genetically inherited. At one time there appeared to be evidence of poverty from studies reporting gesture imitation in newborn infants (Meltzoff & Moore, 1977). However, a range of methodological concerns about this research (e.g. Anisfeld, 1996; Jones, 2006; Ray & Heyes, 2011) have now been vindicated by a large-scale study which, using gold standard methods to test more than 100 infants with 11 gestures, found no evidence whatever that infants could imitate at 1, 3, 6 or 9 weeks of age (Oostenbroek, Suddendorf, Nielsen, Davis, Clark and Slaughter, 2016). As for genetic heritability, a major twin study found that 42% of variance in imitative ability at 2 years of age is related to shared environment, 28% to environmental factors unique to each twin, and only 30% to genetic influence (McEwen, Happe, Bolton, Rijdsdijk, Ronald, Dworzynski & Plomin, 2007). The authors of the twin study concluded that individual differences in imitation ability depend primarily on the amount that an infant has been imitated in the course of his or her development.

All those citations, numbers and percentages – how many infants, how old, how much variance – changed the texture of this essay; it became geekily specific. The medium fits the message. We believe that the broader debate about cultural evolution needs to get geekily specific about cognitive science; to stop relying on plausibility arguments – whether or not they are based on mathematical modelling – and to start testing hypotheses about the nature and origins of cultural learning against the full range of data from comparative and developmental psychology, human experimental psychology, and cognitive neuroscience.

In sum: we admire the case that Henrich makes against High Church Evolutionary Psychology, but we can't understand why he didn't take it further. Surely an analysis that emphasises the power of cultural evolution should consider – if only ultimately to dismiss – the possibility that the power extends to the shaping of cognitive processes. Because SOOS doesn't consider this possibility, or engage with cognitive science more generally, it puts cultural learning at the heart of what makes humans special but does not offer a coherent account of what cultural learning is, or a compelling story about where it comes from.

Evolution

It has recently been suggested that the California school isn't really serious about the 'selection' in their 'cultural group selection' story. That is, their models are kinetic, meaning they are mathematical and involve a plurality of things interacting, but we shouldn't try to interpret them as *Darwinian*, in the way that biologists model natural selection (Lewens 2015). This may be right, but it is somewhat difficult to reconcile with the way Henrich writes in SOOS. He comes across as intending his story, about how humans came to be the way they are, to be an evolutionary story in the standard biological sense – one that invokes a mechanism of blind selection to explain why humans are successful. It is non-standard in that, in addition to the orthodox biological process of natural selection of reproducing organisms, Henrich invokes an additional process (what dual inheritance theorists call a second 'channel') of *cultural* selection. But he seems to suggest that both processes are adaptive – that is, supportive of human success. We won't address the question of what Henrich *really* means here, however. If there is a parallel set of theoretic terms in play then it would be extremely helpful if someone could provide a translation manual to allow communication between

the cultural and standard biological perspectives, and to explain just how the kinetic models are able to explain the evolution of human cultural adaptations if *not* by reference to a Darwinian selection process.

Suffice to say here that we think the cultural selection process *could be* conceived in standard selective terms, and that we *would like* to see the view of cultural selection as a proper selection process developed in sufficient detail that claims about its operation in human history can be evaluated and tested. More specifically, we would like to be able to generate a clear account of the selection models that underlie talk of interacting selection channels; and to understand in exactly what sense we can expect the cultural process to be adaptive.

What are the underlying selection models?

A selection model should be explicit about what are the entities that are reproducing or replicating differentially, and about what it means to call such entities ‘fit’ *beyond* simply acknowledging that some reproduce faster than others. In other words, we should say what the focal units are, what makes some units better than others, and why we should expect the better units to become more populous over time. Viewed in this context, there are at least four different models that might find support in Henrich’s writing:

A) Cultural evolution: Cultural products are selected, by persons, in such a way that their design improves over time, because better-designed products are copied preferentially over poorly designed products.

A cultural product might be an artefact or item of technology; a process or method; or an idea. By ‘design’ we mean the product’s suitability to its immediate function – how well it is built for its purpose. For example, a well-designed hammer is simply good for hammering. There is no implication of intentionality except in so far as a human person will determine what the purpose of the product is. A hammer may be very well-designed, in this sense, even though no person gave its construction any great thought or planning.

We might give process A the catchphrase ‘*Survival of the best-designed cultural products*’. This reconstructed model is based mainly on Henrich’s wonderfully explicit “toy example” given on page 56. Here the objects undergoing selection are successful in so far as they spread – they spread when new tokens come to be possessed by other agents. Thanks to a correlation between the copy rate of such objects and their design adequacy (i.e. how good they are at achieving some local goal) their design improves over time.

B) Culture-driven genetic selection: Genes are selected for their fitness benefits to individual humans, but in an environment that has been shaped by culture.

Process B is standard individual-level natural selection, except that genes are selected in an environment which has been altered by the products of process A. Process B is the default candidate for explaining the evolution of reduced jaw muscles, the dropping of the larynx and other examples given in chapter 5. An appropriate catchphrase here might be ‘*Survival of the fittest individuals, in a cultural environment.*’

Models A and B are each coherent, and describe processes which may work either in concert (as when the cultural selection of dairy farming methods drives the genetic selection of lactose tolerance) or in opposition to one another (as when the cultural selection of cheese production methods impedes the genetic selection of lactose tolerance). However, these models don't seem to accommodate all of the examples Henrich describes, and neither do they fit with his characterisation of the cultural evolution process as group-structured. Here are some further possibilities:

C) Selection of cultures: Whole cultural packages differentially survive, in such a way that cultures become better-designed over time, because worse designed cultural packages attract fewer recruits and become extinct more often.

The catchphrase for this might be '*Survival of the top-recruiting cultural packages*'. It is similar to model A, in that the units bearing fitness are themselves cultural, rather than being made of flesh and blood, but the units are higher-level because their maintenance requires many people interacting. If, in process A, it is the song not the singer that is selected² then we might describe process C as selecting whole symphonies, but not particular orchestras, or particular members of orchestras. In other words, the units in process C are not genetically defined – it doesn't matter which actual humans are practising the relevant cultural behaviours. A cultural practice of this sort may achieve a fitness gain without any new babies being born, if only some new adherents are recruited.

This model would suit the story about pig farmers (p.172), as well as the Dinka/Nuer example that appears in Richerson and Boyd (2005). The package of solicitous pig-care practices does better than the more laissez-faire package, in that a whole tribe of people switches over to the more solicitous package. Model C has got a sort of group element, in that it is a higher-level, or holistic version of channel A. Are the models described in A and C merely different perspectives on a single process? It depends on what exactly a 'cultural package' is, and what it means for one to be fit, successful. If C were merely a more inclusive version of A, we might think that a well-designed cultural package is simply one that incorporates a large number of well-designed products. Sometimes this might fit Henrich's thinking - a successful culture outcompetes others because it has better technologies, better weapons. On the other hand, cultural packages might be understood more holistically, as group-spanning artefacts and ideas such as pig-raising programmes, irrigation systems or religions. The design of such higher-level products can be evaluated in straightforward terms – an irrigation system is good in so far as it distributes water, for example. But this holistic model might also accommodate somewhat murkier higher-level functions – such as group cohesiveness (for religious systems). Unfortunately, Henrich never tells us what it means for a culture to be successful, beyond the fact that they enjoy greater growth and/or persistence than the other cultures, which leaves us only with the tautologous form of 'survival of the fittest'.

Finally, there is another potential model:

D) Group-level genetic selection: Genes are selected for their benefits to group fecundity and/or persistence – helping their group to outcompete rival groups - *rather than* because they are more successful than rival alleles within their own groups.

² To use Doolittle and Booth's appropriate phrase (Doolittle & Booth 2016).

This is standard group or multilevel selection, where we locate selection at the group level if the fitness variance appears predominantly *between* groups, rather than *within* them (Okasha 2006)³. This process could be called '*Survival of the fittest groups*'. Note that if genetic groups have distinctive cultural traits, then it could be the case that the cultural traits are the drivers of group-level genetic selection. In that case, the explanation for the success of the fittest groups would be cultural. The primary difference between processes C and D is that the entities which compete in process D are genetically defined (in other words, a group is a particular lineage of people, regardless of what cultural practice they engage in) while in process C the groups are culturally defined (a group is whoever engages in a particular cultural practice, whichever genetic lineage they belong to). Richerson & Boyd seem to assume something like process D when, in the context of their Tribal Social Instincts Hypothesis, they assume that groups with innate or hard-wired norms of in-group cooperation and out-group hostility die the least in violent conflicts with other groups (Richerson & Boyd 2001). Each ethnolinguistic tribe is assumed to vary from the others, while being internally culturally homogeneous.

Ideally, in discussing particular episodes in human history we would be explicit about which selective model or models we are positing as in play at that time. It is easier to be non-specific and simply gesture at a plurality of candidate models every time, but it is less explanatory. A more serious problem is that the distinct selection processes can interact with one another in various ways, with varying evolutionary consequences. The processes might reinforce one another, or might undermine one another, and might cancel each other out to produce evolutionary stasis. This is worrying because Henrich seems only to consider one of the ways in which they might interact, which is autocatalytically to reinforce one another (p.57) and occasionally seems to be implicitly combining different selection models in ways which are *not* coherent.

One worry is that he is assuming mutual reinforcement between processes C and D, so that the cultural packages selected in process C enhance the biological fitness of their users, such that the genes of those users increase in frequency relative to the genes of users of rival packages. This seems to be the assumption behind one section devoted to considering the question whether group structure has persisted long enough to have elicited a genetic response, "Did intergroup competition shape the social worlds that our genes and psychology faced over the long run during human evolution?" (175).

This is a worry because the groups picked out by process C are not genetically defined. Henrich explicitly expects migration and transmission to act as key mechanisms for the spread of culture (p.168), and yet migration and transmission will confound group-structured genetic selection, because they imply that there simply are no fixed sets of genes for users of particular cultural packages. As Jared Diamond has emphasized, a new language might spread because it sounds musical to an indigenous population's ears, *or* because the invaders slaughter and replace that population (Diamond, 1997, p. 324). In the former case, but not the latter, the genetic identity of the group defined by the language shifts as a whole new group of people is added. What is disturbing is

³ A different way of prioritising group structure would be to invoke a contextualist sense of group selection. Here the group is posited as determining the context in which the genes or artefacts evolve, without requiring variance between groups (Okasha 2006).

that Henrich seems entirely insensitive to the different evolutionary consequences of these mechanisms, simply placing violent conflict, biased migration and prestige-biased copying under the single banner of “mechanisms through which intergroup competition operates” (p. 179). If a group expands by *replacing* another group, then we have a case which fits model D, but if it expands by incorporating immigrants or by converting people from other groups to its ways, then model C fits better. If there is a mixture of both replacement and conversion, as in the case of the Pama-Nyungan spread (p. 176), then only model C will work, and we shouldn’t expect any genetic response to be group-structured.

In general, it is inconsistent to assume that cultural products are often transmitted between groups, *and* that genetic evolution acts at the level of human groups, *if* culture is supposed to be a significant driver of human genetic evolution. Either inter-group transmission and migration play a *smaller* role in determining the number of adherents a culture enjoys than do murder and replacement, or we shouldn’t expect to see much genetic variation at the level of human groups.

In what sense is cultural evolution adaptive?

In addition to the confusion about the nature of the units that are differentially copied in the underlying selection models, there is also a lack of clarity about what it means, in the cultural case, for the units to be successful – what *drives* their differential copying. It is relatively clear what it means for cultural products to be successful⁴. Good hammers hammer well. A well-designed irrigation system distributes a lot of water. A well-designed system of numerals is good for counting with. *The million dollar question is, why should this property – design adequacy – correlate with copy rate?* And what is the role of humans as mediators of the copying process? One possibility is that humans choose preferentially to copy cultural products that they recognise as well-designed, but Henrich excludes this because he wants the process to be blind – for cultural evolution to be smarter than we are. But in that case why, and in what sense, should we expect cultural products to get better over time?

We know what it means to be successful in the genetic models B and D – genes are successful just when humans that bear them have more babies. In order for there to be any causal traffic between cultural and genetic processes of selection, there has to be some kind of connection between the copy rate of cultural products and biological fitness. Yet again, Henrich wants to deny that humans exercise their intentionality to adopt cultural products that they recognise as beneficial for their biological fitness. So, what has cultural evolution got to do with the rate of production of baby humans?

It is not obvious that Henrich really confronts this problem of why we should expect to see the accumulation of those cultural products that are good for human fitness, rather than of those that are pleasant to humans but bad for their fitness, or of those that are neutral in respect of human fitness but good for the replication rate of cultural products. There follows a suggestion about how Henrich *could* develop his account so that cultural fitness is tied to biological fitness. We advance it, not as an interpretation of what Henrich himself means to say, or even as our own view, but as an

⁴ Henrich 2004 talks specifically about the property of ‘skilfulness’ – how much learning must be done in order to use a tool – as increasing under cultural evolution. But it seems paradoxical to suggest that this is the property that is selected *for*. Why would some artefact be copied more in virtue of being harder to use?

illustration of the kind of conceptual work that we believe needs to be done to examine the possibility that cultural evolution is based on selection.

Henrich devotes a considerable amount of space to the treatment of learning biases, saying that we *choose* models to learn from, and that this gives rise to cumulative improvements in cultural products. But he never gets specific about the role of this “second-order cultural learning” (p. 119) in supporting cumulative cultural evolution. The biases may simply be another feature of Boyd and Richerson’s account which Henrich has inherited. Yet maybe the social learning biases can maintain the all-important connection between the copy rate and the design value of cultural products and/or packages. It needs to be the case that variant artefacts or ideas that constitute improvements—that are better suited to the purposes for which the human user is employing them – are copied more often than are novel variants that constitute deteriorations in design. Improvements must get copied more often than do errors. Henrich doesn’t want to have to assume that any agents are smart enough to actually spot the difference between good and bad designs. But perhaps his agents are biased in their copying behaviour in just such a way that clever alterations get taken up more often than do dumb alterations, without anyone being clever. In other words, maybe learning biases could lead humans preferentially, albeit *accidentally*, to copy artefacts that are well-designed. For example, prestige bias might lead us preferentially to copy well-designed tools, *if it is the case* that prestigious individuals are more likely to use well-designed tools than they are badly-designed tools. Age-bias might lead us preferentially to copy well-designed tools *if it is the case* that older individuals are more likely to use better-designed tools. Then the necessary correlation between copy rate and design follows. Henrich himself identifies examples in which the correlation between copy rate and design-adequacy fails, such as in the advertising of insurance by basketball stars. Yet learning biases could be enough to make a process of CCE progressive as long as there is *some* positive correlation between those we preferentially copy and the design-adequacy of the conceptual and concrete artefacts that they employ.

Note that even if other animals were to become really high-fidelity imitators who were strongly motivated to copy and so on, and even if they lived in huge groups so that there was plenty of cultural capital available, they’re not going to get anywhere if they copy bad variations just as often as good variations. As long as there is no causal link between design value and copy rate, the only thing such a system will produce is objects that are optimised for getting copied: a memetic process. None of the cumulative ‘improvements’ that turn up would be for the benefit of the agents that do the copying. So, on our suggestion, the human innovation is to correlate good design to copy rate, not by learning to recognise good design, but simply by copying preferentially from agents who happen to be more likely to be using the best tools.

Just how plausible is it to suppose that learning biases could tweak a blind selection process so that the things copied get better suited to human goals, instead of simply better at getting themselves copied? It might be plausible that copying from older individuals will mean copying better-designed tools, if we assume that older individuals have had more time in which to use trial-and-error individual learning to improve their tools. The trouble is that the source of the improvement here seems to be the individual learning, more than or as much as any blind cultural selection process. Similarly, success-bias and skill-bias might be strong candidates for tying copy rate to good design, but it’s not easy to explain how these biases could work in such a way that excludes smartness. What about prestige bias – is it plausible that copying from the prestigious will mean copying well-

designed artefacts, more often than not? This depends on the details of how and why individuals come to be prestigious. If generosity is an important cause of prestige (p.128) this may make prestige bias a good candidate for tying copy rate to good design in respect of the sorts of higher-level cultural packages, such as religion, which function to enhance group cohesion or cooperation. At least Henrich makes a good case for prestige being a property that we can track fairly unintelligently. Put simply, if we blindly preferentially copy prestigious individuals, and prestigious individuals tend to be more generous than others, then we might expect everyone to become more generous as a result. However, it is harder to see how this could work to improve the design of tools with non-social functions.

Finally, if we are to keep the idea that cultural evolution can feed back into genetic evolution, then design adequacy must be connected to biological fitness in some way. To see that such a connection cannot be taken for granted, consider the contraceptive pill, which may be very well suited to its purpose, but which reduces biological fitness. Human ends do not always coincide with biological ends. Henrich documents how lots of people flocked to copy an eminent climber of Mount Everest.....who went on to freeze to death (p. 81). But again, it's enough if we go along with the idea that there has been *some* positive correlation between using effective tools and having lots of children, at least for most of our human evolutionary history. One advantage of positing a less than perfect connection is it leaves the door open, again, to a maximand that isn't reducible to *individual* biological fitness: some sort of social or group fitness, for example. But clearly there is a huge amount of further clarification still left to do.

Maybe these concerns are inappropriate because, as Lewens suggested, we are not supposed to view the selection processes at work in Henrich's account of cultural evolution as Darwinian. But if someone *does* want to understand cultural selection as a fully Darwinian, *adaptive* process, then these *cui bono* details (for whose benefit? Dennett 2001) cannot be avoided. In particular, there needs to be some clarification of how humans have secured the preferential copying of the most useful cultural variants without being smart.

Conclusion

SOOS tackles an important question – why a bald, wimpy species of primate has come to enjoy such global dominance – in an engaging way, drawing on multiple disciplinary resources to construct explanations for puzzling phenomena, such as little quirks of human language systems (Chapter 13), and why humans but no other primates might have come to value cultural capital (Chapter 16). However, there is a lack of clarity about the nature and origins of cultural learning, about the models of selection which underlie the account, and about the nature of the improvements that are supposed to result from cultural selection processes. Rarely will a trade book satisfy an academic readership but, from a stylistic perspective, some of these problems might have been solved by giving an overview of principal claims, and an outline of the book's structure, in one of the introductory chapters. Nonetheless, even if they have to work hard to uncover all the arguments, academic readers can expect to find plenty that is provocative and interesting in "The Secret of our Success".

Henrich is very far from being alone among cultural evolutionary theorists in glossing over important questions about cognition and selection. This kind of idealisation was adaptive when research on cultural evolution was getting off the ground (Boyd & Richerson, 1985), but the strategy is rapidly outliving its usefulness. Now that serious, interesting, important, broad-brush work on cultural evolution has been in progress for more than 30 years, it's time to get down to the empirical and conceptual details.

References

- Anisfeld, M. (1996). Only tongue protrusion modeling is matched by neonates. *Developmental Review*, 16, 149–161.
- Chomsky, N. (1975). *Reflections on language*. New York: Pantheon Books.
- Dennett, D. C. (1978). Three kinds of intentional psychology. *Perspectives in the Philosophy of Language: A Concise Anthology*, 163-186.
- Dennett, D.C. (2001). The evolution of culture. *The Monist*, 84(3), 305-24.
- Diamond, J. (1997) *Germ, guns and steel. A short history of everybody of the last 13 000 years*. New York: Norton.
- Doolittle, F. & Booth, A. (Forthcoming) It's the song, not the singer: an exploration of holobiosis and evolutionary theory. *Biology and Philosophy*.
- Fogarty L, Rendell L, Laland K.N. (2012) Mental time travel, memory and the social learning strategies tournament. *Learning & Motivation*, 43, 241–246.
- Frankish, K., & Ramsey, W. (2012). *The Cambridge Handbook of Cognitive Science*. Cambridge University Press.
- Galef, B. G. (1971) Social effects in the weaning of domestic rat pups. *Journal of Comparative and Physiological Psychology*, 75, 341-357.
- Henrich, J. (2016) *The secret of our success: how culture is driving human evolution, domesticating our species, and making us smarter*. Princeton: Princeton University Press.
- Heyes, C. M. (2012) Grist and mills: on the cultural origins of cultural learning. *Philosophical Transactions of the Royal Society B*, 367, 2181–2191.
- Heyes, C. M. (2013) What can imitation do for cooperation? In K. Sterelny, R. Joyce, B. Calcott, & B. Fraser (Eds.) *Cooperation and Its Evolution*. MIT Press.
- Heyes, C. M. (2016a) Blackboxing: social learning strategies and cultural evolution. *Philosophical Transactions of the Royal Society: B*, 371, 20150369.

- Heyes, C. M. (2016b) Who knows? Metacognitive social learning strategies. *Trends in Cognitive Sciences*, 20, 204-213.
- Heyes, C. M. (in press) When does social learning become cultural learning? *Developmental Science*.
- Heyes, C. M. (in prep) *Cognitive gadgets: the cultural evolution of thinking*.
- Heyes, C. M. & Pearce, J. M. (2015) Not-so-social learning strategies. *Proceedings of the Royal Society of London: B*, 282, 20141709.
- Hoppitt W & Laland KN. (2013) *Social learning: an introduction to mechanisms, methods, and models*. Princeton, NJ: Princeton University Press.
- Jones, S.S. (2006). Exploration or imitation? The effect of music on 4-week-old infants' tongue protrusions. *Infant Behavior & Development*, 29, 126–130.
- Laland, K.N. (2004) Social learning strategies. *Animal Learning and Behavior*, 32, 4–14.
- Laland, K.N., Odling-Smee, J., Hoppitt, W. & Uller, T. (2013) More on how and why: cause and effect in biology revisited. *Biology and Philosophy*, 28(5), 719-45.
- Lewens, T. (2015) *Cultural evolution: Conceptual challenges*. Oxford: Oxford University Press.
- McEwen, F., Happe, F., Bolton, P., Rijdsdijk, F., Ronald, A., Dworzynski, K., & Plomin, R. (2007). Origins of individual differences in imitation: links with language, pretend play, and socially insightful behavior in two-year-old twins. *Child Development*, 78, 474–492.
- Meltzoff, A.N., and Moore, M.K. (1977). Imitation of facial and manual gestures by human neonates. *Science*, 198, 75–78.
- Okasha, S. (2006) *Evolution and the Levels of Selection*. Oxford University Press.
- Oostenbroek, J., Suddendorf, T., Nielsen, M., Redshaw, J., Kennedy-Costantini, S., Davis, J. Clark, C. & Slaughter, V. (2016). Comprehensive longitudinal study challenges the existence of neonatal imitation in humans. *Current Biology*, 26, 1334-1338.
- Ray, E., and Heyes, C. M. (2011). Imitation in infancy: the wealth of the stimulus. *Developmental Science*, 14, 92–105.
- Richerson, P.J. & Boyd, R. (2005) *Not by genes alone: How culture transformed human evolution*. Chicago,: University of Chicago Press.
- Richerson, P.J. & Boyd, R. (2001) The evolution of subjective commitment to groups: A tribal instincts hypothesis." *Evolution and the Capacity for Commitment* 3 : 186-220.
- Shallice, T., & Cooper, R. (2011). *The Organisation of Mind*. Oxford University Press.
- Tomasello, M. (1999) *The cultural origins of human cognition*. Cambridge, MA: Harvard University Press.

Tomasello, M. (2014) *A Natural History of Human Thinking*. Cambridge, MA: Harvard University Press.