Hanging On To The Edges:

What is cultural evolution like?

Different problems require different tools
- Tim Lewens, *Cultural Evolution: Conceptual Challenges*

For over twenty years, I have been having an on again/off again affair. The other party in the affair is appealing and effortlessly elegant. She promises to wash away the mundane, equivocal, hesitant, mutually contradictory strata of my work as a behavioural scientist and replace them with something simpler, crystalline and more powerful. I encounter her from time to time in the course of my professional duties. We write a paper or two together. I expect summits at Camp David. But within a year or two I start to pull away. She is brittle behind the mask. I start to worry about how the bills will really get paid. When the mirrors have stopped dazzling and the smoke has cleared, I realise I still have all the problems I had before.

The other party in the relationship is an idea. Actually, two linked ideas: (i) that cultural change is a Darwinian process; and (ii) that because (i) is true, social science can be substantially simplified under the rubric of a single body of theory that does the same job, in the same way, as evolutionary theory does for genetic evolution. These ideas have been knocking around for about forty years. They have their passionate adherents[1]. But they continue to attract scepticism, and despite all the conceptual discussion, I don't notice journals of sociology, politics, social anthropology, history, cultural studies and so on being full of empirical calculations of cultural fitness, cultural relatedness, cultural heritability and so forth, in the way that journals of evolutionary biology are full of the genetic versions of these notions. So what is going on?

There are a number of possibilities. One is that ideas (i) and (ii) are fundamentally correct, and Rome wasn't built in a day. The history of science shows us that the right idea takes a long time to rise up through layers of inertia, tradition and disciplinary resistance. Plate tectonics, for example, took about 60 years from first, derided claims to universal acceptance. The acceptance took the form of a characteristic S-curve: very slowly rising for a long time, then a phase of rapid spread, then the slow mopping up of the few remaining non-believers. So as a cultural Darwinian, you must tell yourself that

---

[1] Notable adherents are: Dawkins, R. (1976). *The Selfish* Gene (Oxford: Oxford University Press); Mesoudi, A., A. Whiten and K.N. Laland (2006). Toward a unified science of cultural evolution. *Behavioral and Brain Sciences* 29: 329-83; Mesoudi, A. (2015). Cultural evolution: a review of theory, findings and controversies. *Evolutionary Biology* 43: 481-97. Before I annoy anyone any more than I need to, I should distinguish between *cultural evolutionary thinking*, and *cultural Darwinism*. The former is simply the attempt to understand the population-level consequences over time, for human societies and their cultural attributes, of individual patterns of learning and cognition. It is broader than *cultural Darwinism*, a subset of cultural evolutionary thinking which sees cultural change as a process of Darwinian selection and hopes through that insight to radically transform the social sciences. Cultural evolutionists are not necessarily committed to cultural Darwinism, and indeed some have vociferously criticised it. See Lewens, T. (2017). *Cultural Evolution: Conceptual Challenges* (Oxford: Oxford University Press) for discussion.

you are just entering the accelerating phase on the S-curve; *this year*. My problem is: I thought that 20 years ago, when I wrote my first cultural Darwinian papers[2]. I am still waiting.

Another possibility is that idea (i) is wrong, and hence idea (ii) also fails, but there are other reasons people cling to them. Research in the humanities and social sciences is in slow decline (not without a fight). Eighty years ago, to a fairly reasonable approximation, the humanities and social sciences were what universities did. Today, also to an approximation, universities have a dual role: they teach students in humanities and social sciences, and they do research in biology. A glance at the difference in teaching load, and research and infrastructure funding, between my faculty in my university (Biomedicine), and the Faculty of Humanities and Social Sciences, is instructive. We are researchers with expensive labs and technical support teams who give the odd lecture; they are teachers who occasionally manage to scrape the time to write articles and books. Anything that seems to offer the humanities and social sciences the possibility of getting what biology has had seems worth grasping at. So idea (i) does not quite die. Sadly, those who want to save the humanities and social sciences through a Darwinian theory of culture are probably looking in the wrong place. The recent growth of biology is almost entirely in cellular and molecular work, a part of biology largely free from the guiding light of Darwin's dangerous idea. The driving forces have been rapid technical progress in what researchers can measure, and the computational firepower to mine the resulting big data. So that's where a lot of the smart money in social science is going to go too.

As so often in life, I find myself somewhere in the middle ground. The analogy between genetic and cultural evolution is strong enough that it continues to capture my theoretical attention[3]. On the other hand, it's not straightforward enough for ideas (i) and (ii) to get off the ground in a major way. I don't expect the revolution imminently. Hence my on/off affair. Hanging on to both edges, as usual.

§

Genetic and cultural evolution are not exactly isomorphic. Everyone admits that. On the other hand, there are some general similarities: something gets transmitted from individual to individual; some things spread and others go extinct; there is a kind of descent with modification, and so on. So the issue is: what do we do with this partial similarity? We could either: define Darwinian processes rather narrowly, and thereby include genes but exclude culture; or find broader ways of defining Darwinian processes, so as to include the cultural case as well as the genetic[4]. So clearly the answer we get to the question of whether cultural change is Darwinian will depend on the definition of 'Darwinian' we adopt. A more fruitful avenue, to my mind, is to ask: what special job does evolutionary theory do for organismal biology, and what are the properties of genetic evolution in virtue of which it can do that job? Then we can assess the extent to which cultural evolution has those properties, and hence whether a Darwinian 'cultural evolutionary theory' could do that job. I'm in an off-again phase, and so I am going to conclude that cultural evolution generally lacks the properties and hence 'cultural evolutionary theory' (thought of in this particular way) can't really do the job. This much has often been said before, by better people than me. Perhaps I have a slightly more unusual insight, though, which is that the real problem for the hope of a unified Darwinian 'cultural evolutionary theory' is that different cultural cases are very different from one another, and hence approximate the genetic situation to different degrees. This is a serious blow to hope (ii).

---

[2] For example Nettle, D. (1999). Functionalism and its difficulties in linguistics and biology. In *Functionalism and Formalism in Linguistics* (eds. M. Darnell et al.), Amsterdam: Benjamins, pp. 445-462.

[3] Most recently in El Mouden, C., J-B André, O. Morin and D. Nettle (2014). Cultural transmission and the evolution of human behaviour: A general approach based on the Price equation. *Journal of Evolutionary Biology* 27: 231-41.

[4] See Claidière, N., Scott-Phillips, T. C., & Sperber, D. (2014). How Darwinian is cultural evolution? *Proceedings of the Royal Society B: Biological Sciences* 369: 20130368.

It's a commonplace that you have Darwinian evolution whenever you have variation (different individuals in a population have different traits); heredity (offspring resemble their parents); and differential reproductive success (the descendant generation differentially samples from the ancestral one, or equivalently, different individuals have different chances of becoming ancestors). This much is true, but up I think we need to build the requirements up more slowly. First, there must be a clearly defined population of *individuals* through time; you need to know what your individuals are. Second, within that population, you must be able to identify which individuals are *descendants* of which others, and which are not. Without being able to do this, there is no hope of measuring reproductive success, since the very notion depends on descendant-counting. Third, these individuals need to have *traits*: characters, discrete or continuous, that you can measure, and hence characterize straightforwardly the extent to which descendant is like ancestor.

With these requirements in place, we can characterize the way any particular trait changes from one generation to the next. This was famously done by George Price, in the Price equation[5]. So general and important is this equation that, to my knowledge uniquely for an equation, it has a movie based on it, *wΔz* (directed by Tom Shankland, 2007; it's a horror story). The Price equation says, in words, that in each generation:

Total evolutionary change in the trait =

A bit due to selection +

A bit due to average transformation

The Price equation also tells us how to compute the value of each bit. The bit due to selection is exactly the covariance between the value of trait, and *fitness*, where fitness is the individual's number of descendants in the next generation, divided by the population average number of descendants. A covariance is like a correlation: it can be positive, negative or zero. So let's say that the trait is nose length. If it is the case that the longer your nose, the higher your reproductive success on average, then the covariance between nose length and fitness is positive, and the value of the bit due to selection is positive. This means selection is making noses longer from generation to generation. If longer noses tend to be associated with reduced reproductive success, then the covariance of nose length with fitness is negative, and hence selection is making noses shorter. And of course, the length of your nose may have no systematic relationship with reproductive success, in which case, the bit due to selection has a value of zero, and there is no directional selection on the trait.

Then there is the bit due to average transformation. Imagine a case where, because of some strange quirk of genetics or development, offspring always had noses that were a bit longer than the average of the lengths of their two parents' noses. It's easy to see that noses would get longer over evolutionary time, even in the absence of any natural selection. In fact, they could get longer over evolutionary time even with some natural selection acting in the opposite direction. The Price equation tells us exactly when this will happen: when the average amount by which an offspring's nose length exceeds those of its parents (the bit due to average transformation) exceeds the negative covariance between nose length and fitness (the bit due to selection). This is because, to get the total evolutionary change from the Price equation, you simply add the two bits on the right-hand side together.

It's important not to confuse random mutation or imperfect heredity with average transformation. Let's say there is quite a lot of genetic mutation, so that offspring nose length is not perfectly predicted by parent nose length. If an offspring's nose is just as likely to be a bit shorter than their parents' as a bit longer than their parents', then *on average*, their nose is neither longer nor shorter, and so the bit due to average transformation is still zero. The genetic mutation averages itself out as it were, and the

---

[5] Price, G.R. (1970). Selection and covariance. *Nature* 227: 520–521.

total evolutionary change comes from the bit due to selection alone, even though individuals do not perfectly resemble their parents.

§

With this exposition in the bank, we can begin to ask what properties genetic evolution has that allow it to do a special job for organismal biology. And it does do a rather special job. A relatively straightforward mode of theorising, in which trait evolution is explained on the basis of higher or lower fitness, can be brought to bear in much the same way on any biological trait, be it the dimensions of the hummingbird, the shape of fish eyes, the propensity to help others reproduce, or the way animals forage. The mode of theorising can be used regardless of what the trait in question actually is, and in particular, without knowing anything about the details of the molecular genetic mechanisms involved. Why can we do this?

The first reason we are able to do this in the genetic case is that the relations of ancestry and descent are straightforward. I have just two genetic parents. No one else has influenced the length of my nose (heritably, that is). Those same two individuals are my parents in respect of all of my other traits, not just nose length. The ancestor-descendant link points in one direction only: I can't back-influence the heritable traits of my genetic parents. And how many parents I have does not depend on the lengths of their noses. That sounds bizarre, but is not guaranteed in the cultural case. For example, I might sample the way of life of the first few people I encounter. If it seems to suit me, fine, I follow it, but if it seems dreary, I might go looking for other people to copy. This is me shopping for cultural ancestors on the basis of the traits they offer, something we don't get to do with our genes. Because of the straightforwardness of the ancestor-descendant mapping in the genetic case, you can readily count offspring and measure reproductive success. And then it's easy to compute the value of the bit of evolutionary change due to selection: measure the trait you are interested in, count descendants, apply the formula for a covariance.

The second reason fitness is central to organismal evolution is that for biological traits, the bit due to average transformation is usually zero or negligible. The conclusive evidence for this is that when, in *wΔz*, the evil serial killer carves the Price equation onto the flesh of his victims, mindful of parsimony and apparently having the biological case in mind, *he only includes the term for the bit due to selection on the right-hand side*. He can simply leave out the term for the bit due to average transformation because, usually, genetic reproduction systems have no particular transformational drive one way or the other. Thus the Price equation, for genetic cases, usually reduces to evolutionary change being equal to the covariation between the trait value and fitness.

Why is average transmission negligible as a source of change in genetic systems? It is because, in some profound sense, the function of DNA replication mechanisms is to indifferently reproduce whatever is thrown at them. That's their job. DNA replicase is indifferent between replicating a cytosine and replicating a guanine, indifferent indeed to what if anything the particular stretch of DNA it is currently copying actually does. It simply *has no interests* other than to fulfil its evolved role of making DNA into more DNA. To the extent that mutation happens (and it does, though overall fidelity is high), this mutation can be fairly analogised, as it often is, to 'mistakes' or 'imperfections' in the replicative process. And at reproduction, fair meiosis generally ensures that no variant gets a leg up, on average, over any other.

The profound content-indifference and impartiality of DNA replication provides us, as scientists, with the option of abstracting away from a lot of the details of how replication and reproduction actually work in each particular case. You can usefully treat genes in populations as if they were simple beans being drawn from a bag[6]. To think about the evolution of nose length, you don't really need to know about the molecular details of which stretches of DNA influence nose length and how, at least in the

---

[6] Haldane, J.B.S. (1964) A defense of beanbag genetics. *Perspectives in Biology and Medicine* 7:343–59.

first instance. This is because you can take it for granted that those molecular details, however they work out, come down to offspring of long-nosed parents having long noses, plus some effectively random noise. So we can make a great deal of progress just by knowing that nose length is heritable, and measuring how it correlates with fitness, without getting bogged down in the messy biology of the specific case. Indeed, the messy biology of the mechanisms in each particular case is still largely unknown to us: a black box which we have scarcely as yet peered inside. And because in each biological case, we have been able to abstract from the details of the molecular and developmental processes involved, then all normal biological cases are effectively like one another from an evolutionary theoretical point of view. Transmission is by fair replication from one or two parents, and selection (plus drift) are, to a first approximation, the drivers; the rest is about the impact of the trait on fitness in populations. The general formulae apply. This is the sense in which we can sensibly talk of *an* evolutionary theory, rather than one theory for the evolution of eye shape, based on the developmental biology of eyes, one theory for the evolution of blood proteins, based on the physiology of blood proteins, and so forth.

§

Now let's start to think about cultural cases. We start where I started my career, with the evolution of words[7]. For example, consider Zipf's law, which states that words that are used more frequently in a language tend to be shorter, whilst rare words are longer. This looks pretty much like an adaptation: it benefits speakers in terms of overall articulatory effort if the shortest available word forms are used up on the meanings we need most often, and longer words forms saved for meanings we don't need to utter very often. And it's tempting to characterise the process producing the pattern in terms of selection. Doing so, in fact, goes right back to Darwin, who noted in *The Descent of Man*:

> As Max Muller…has well remarked:-"A struggle for life is constantly going on amongst the words and grammatical forms in each language. The better, the shorter, the easier forms are constantly gaining the upper hand, and they owe their success to their own inherent virtue."[8]

Let's unpack how this might work. People have a set of cultural parents from whom they learn their word forms. We can assume that they learn their word forms early in life, and the forms are fixed thereafter. So that means word-form learning really is a bit like genetic inheritance, with ancestor-descendant arrows pointing in one direction only, from older people to younger. The set of cultural parents is broader than just their genetic parents of course, but the Price equation can be generalised to an arbitrary number of parents, and indeed can be generalised to a case where parenthood is a matter of degree: you have many cultural parents, and some are more influential than others[9]. Fitness then becomes not the number of descendants you have, but the average strength of your influence on all the individuals in the next generation. No intrinsic problem there; though, given that the set of people you learn from can be different for different traits (I learn playground games from my peers, science from my teachers), it does follow that every individual has not just one cultural fitness, but indefinitely many cultural fitnesses, one for each cultural trait.

The main quibble with analogising the emergence of Zipf's law to adaptation through natural selection is that it is not clear the adaptation arises from the selection bit of the Price equation, rather than the average transmission bit. It could be that as you grow up, you hear various idiosyncratic variant word forms spoken around you, and you have a bias towards adopting in your own speech those variants that give you short word forms for frequent meanings. That would be a kind of selection. But an alternative (not mutually exclusive) mechanism is that as you use language, you tend to spontaneously

---

[7] Nettle, D. (1995). Segmental inventory size, word length, and communicative efficiency. *Linguistics* 33: 359-67.

[8] Darwin, C. (1871). *The Descent of Man, and Selection in Relation to Sex.* London: John Murray, pp. 465-6.

[9] See El Mouden, C., J-B André, O. Morin and D. Nettle (2014). Cultural transmission and the evolution of human behaviour: A general approach based on the Price equation. *Journal of Evolutionary Biology* 27: 231-41.

shorten words or phrases you utter frequently. You might do this even if you have not heard the people you learn from do so. Geddit? This could be the source of Zipf's law. Word forms tend to start life long, and if they are used often, speakers spontaneously contract them through economy of effort. If that's right, then the adaptive change is not actually due to *selection*, but rather a particular bias in average *transmission* (a person's habitual word form for a common meaning will be a bit shorter than the average length of that word form in the models from whom that speaker learned).

Now whether word evolution is due to transformation or selection is, you might say, a rather unimportant technicality. Perhaps. But there have been very cool experiments in recent years using so-called *transmission chain* or *iterated learning* paradigms. Here, one participant learns something (a starting stimulus furnished by the experimenters); a second participant learns from the first; a third participant learns from the second; and so forth. These experiments have been applied to cultural content as varied as stories, communication conventions, rhythmic patterns, statistical relationships, and many other things[10]. They show in fascinating detail how cultural change can be rather fast, and decidedly non-random: typically, across just half a dozen links of the chain, the learned content changes in patterned ways. The contents go from being essentially random in the first generation to having the rich structure seen in real cultural representations: communication systems become grammatically regular; rhythms acquire the regular pattern of strong and weak beats you see in all music; and complex random statistical scatters are reduced to simple, memorable stereotyped relationships.

I love these experiments. They show us how culture evolves. But there is no selection going on. There can be none by design, since every participant has exactly one cultural ancestor and exactly one cultural descendant. The bit due to selection in the Price equation is therefore exactly zero. The rich, non-random shaping of the content that we can see in these experiments must be entirely due to average transmission—the way that human participants actively shape the information they are exposed to, in accordance with their purposes, strategies and biases—and not at all due to selection. That's a key difference from genetic evolution. And it raises the question: if you can get good experimental analogues of cultural change in an experimental set-up that excludes any selection, to what extent is selection needed to do the explanatory work for cultural evolution in the wild?

§

They say hard cases make bad law. But there ought to be a similar adage for easy cases and bad generalizations. By choosing the cultural cases that *best* fit the analogy between genetic evolution and cultural change, we overestimate how good the analogy is overall. In the example of word forms considered until now, there are some rather atypical circumstances that obtain. It is reasonable to assume that one learns one's word forms early in life, and hangs on to them thereafter. So the ancestor-descendant links go in one direction only, from older to younger people. But this is not the general case. Consider the cultural evolution of scientific ideas. I have learned a lot about science from my long-term collaborator Melissa Bateson. But I have also transformed what I have learned from her, using my own particular cognitive operations, and back-influenced her in turn. And then of course she has reflected on and transformed those ideas still more, influencing me again. So who is ancestor and who descendant? How will we deal with this if we wish to maintain some parallel between genetic evolution and cultural change?

---

[10] Kirby S., H. Cornish and K. Smith (2008). Cumulative cultural evolution in the laboratory: An experimental approach to the origins of structure in human language. *Proceedings of the National Academy of Sciences of the USA* 105:10681–10686; Ravignani, A, T. Delgado and S. Kirby (2017). Musical evolution in the lab exhibits rhythmic universals. *Nature Human Behaviour* 1: 0007; Griffiths, T. L., Kalish, M. L., & Lewandowsky, S. (2008). Theoretical and empirical evidence for the impact of inductive biases on cultural evolution. *Philosophical Transactions of the Royal Society B: Biological Sciences* 363: 3503–3514.

We could say that a person becomes a new individual in respect of any particular cultural trait every time they have a change of idea. So Melissa Bateson of 2017 is a cultural descendant of, among others, Daniel Nettle of 2016, in the domain of scientific ideas. But Daniel Nettle of 2016 is a cultural descendant of Melissa Bateson of 2015. So we would have to admit that Melissa is one of her own cultural grandparents in the domain of scientific ideas. Maybe that's ok. But it means that, for culture, Melissa does not just have indefinitely many different fitnesses. She is also indefinitely many individuals (some of whom are ancestors to some of the others in some domains). And if that weren't complex enough, many of those individuals are alive at the same time. She wrote articles twenty years ago that are still influencing biological ideas in January 2017 via a route other than what she believes in December 2016 (she may even have forgotten what they say). I hope it's clear that the moment you have the possibility of repeated and continuous learning over the life-course (let alone literacy), the whole question of ancestry, descent, and fitness becomes really rather difficult to track; and culture has potential dynamics that cannot be captured with models inspired by beanbag genetics[11].

§

Let's take another example close to my heart: the cultural evolution of my local cross-country running league. If cultural Darwinism can do anything, I really want it help me here, because this venerable institution shows clear evidence of descent with modification, plus evidence of adaptation to human purposes, over the course of its 120-year history[12]. Like other Victorian running leagues, it began with 'hare and hounds' pursuit races. (This is the origin of the name 'harriers' in the names of many running clubs in the English-speaking world; and also the source of the term 'steeplechase': routes were not fixed, but the hares would head for visible landmarks across the countryside, jumping streams and fences.) The runners in the pursuing pack were not racing against each other for most of the run. They would deliberately stay together in a peloton until a final competitive sprint for around the last mile, whose sudden onset was orchestrated by a special runner called the Whipper-In. After 1950, this system was simplified to 'all out' racing: the participants all competed against one another from gun to tape, not just for the final sprint; the hares and Whipper-In were abandoned.

A subsequent innovation was a handicap system, whereby runners who finish high up in one race must start with a handicap in subsequent races of the season. This keeps things challenging for the fastest people, whilst providing some hope for the slower ones. The handicap system applies only to league matches—there is also one cup match a year, where everyone goes off together. The rules for promotion to a higher handicap in league matches are not symmetrical with those for demotion: you can increase your handicap after every match, but only decrease it at the end of the season. This means that the first group to set off gets smaller and smaller as the season goes on, as more and more people attract handicaps. Thus, people who would have no hope of winning in the first match of the season do so in the last. Then there are complex rules for how the team's score arises from the positions of individual runners (it's a team sport); and, a very recent innovation, there are several divisions for teams, with rules for team promotion and demotion from their division.

I could go on, but I hope you can see that there is an exquisite cultural order here that has emerged by the slow reshaping of institutional tradition to conform to human needs and desires. Is there any way we can fit this slow reshaping into a Darwinian framework? It's not clear that there is variation and selection, since there has only ever been one cross-country league in Northeast England. So there aren't really competing variants with higher or lower fitness. After all the whole point of an institution is that at any given time, all individuals have to sign up to exactly the same rules. I suppose we could say that in some sense the rules of cross-country are in competition with other things people could

---

[11] Strimling, P., M. Enquist and K. Eriksson, K. (2009). Repeated learning makes cultural evolution unique. *Proceedings of the National Academy of Sciences of the USA* 106: 13870-13874.

[12] The authoritative source is Jenkins, A. (2016). *Whipper-In: The Northumberland and Durham Paperchase League, the Early Years, the Forerunner of the North East Harrier League.* Alnwick: Wanney Books.

spend their Saturday afternoons doing; perhaps the rules of cross-country have a fitness relative to golf, say, or gardening. But this is rather different from the genetic evolution of nose length, where noses of a given length are in competition with simultaneously-present noses that are slightly shorter or longer, driving nose length up an adaptive gradient of nose length.

And then there is how change works. There's modification, but is it usefully thought of as inheritance plus mutation? There is a league committee, who are bound by a written constitution. The committee sits down at the end of the season, looks at what went well and what people complained about, and uses reason and argumentation to decide whether there is anything they want to change. So you could perhaps say that there is a kind of virtual variation and selection process, whereby the committee simulate in their minds various alternative possibilities, then choose the one that seems best for the sport. That would be rather evolution-like in way, with the exceptions that the variants with lower fitness never get to actually exist outside of committee meetings.

Given that there is only one set of rules, the ancestor for each rule in each season seems to be the corresponding rule in the previous season—faithful inheritance, with agreed rule changes playing the role of mutation. But maybe that's not right. For example, the recently-introduced divisional system for teams, with its rules for promotion and demotion, is clearly similar to the rules that do the same job for the Football League. So perhaps the cultural trait we should be thinking about is not 'the rule system for cross-country running' but 'divisional systems in competitive sports'. Then we could talk of the fitness of divisional systems being high, as they have colonized new sports like cross-country running. But there might be things in some sports which are *a bit like* divisional systems, fashioned somewhat in the style of such systems, but not exactly the same as them. How would these contribute to the cultural fitness of the people who advocated a divisional system for the Football League? Would they count towards it, or not?

There are harder cases still. I fully expect within a few years, there will be pressure for women and men to run the same course (currently, senior women run two laps and senior men three). If this happens it will reflect the broader social concern about gender equity that we currently see, for example, in the debate about wages in the UK. Harriers bring that broader cognitive and political framing into their leisure activities. If gender-equalization happens, then we are effectively saying that the value of a trait in cross-country running (course length) has been influenced by the existing trait-value, plus a general concern about gender equity that is otherwise manifest in entirely different domains of human activity. This has no parallel in genetic evolution. It is as if you said: the two heritable influences on the length of my nose are the length of my father's nose, and my mother's sense of humour. But culture is pervasively like this. Cognition ranges promiscuously across domains and activities, seeking partial resemblances and relevant reasons, recombining, tidying up, reconstructing one thing in the light of another. This makes even identifying what the traits under cultural evolution are, as well as delineating who is culturally ancestral to whom, very complex at best[13].

§

It is time to limp, bloodied but still determined, towards some more general conclusions. For genetic systems, transmission is achieved by content-indifferent replication mechanisms; and the relationships of ancestor and descendant are straightforward in all cases, and the same for all traits. This allows us bypass the details of how transmission is actually achieved, and go straight to a special type of explanation in which fitness, and the relationship of different organismal designs to fitness, do all the explanatory work. Evolutionary biologists have special names for the thinking that underlies

---

[13] Claidière and colleagues provide a way of addressing some of this through their notion of hetero-impact: a feature in one cultural generation can have a causal impact on a different feature in the next. See Claidière, N., Scott-Phillips, T. C., & Sperber, D. (2014). How Darwinian is cultural evolution? *Proceedings of the Royal Society B: Biological Sciences* 369: 20130368.

this special type of explanation: beanbag genetics (idealizing genes to understand their dynamics in populations); ultimate reasoning (thinking about fitness consequences of a structure rather than the mechanisms that make it); the phenotypic gambit (ignoring the genetic architecture of the traits under study); and the behavioural gambit (assuming brains can deliver whatever is good for fitness without worrying how they might do so). And because you can do these special types of thinking for all traits, there is a coherent sense in which there is *one* evolutionary theory: one tightly-integrated, portable system of tools for working out change in any genetic system. The setting aside of the details of how transmission actually works is only a provisional strategy; and the worry is often voiced within biology that you can't entirely get away with it. In the end, the details of the available mechanisms are probably going to matter for what happens[14]. Still, people have been able to do a lot by making the idealizations and thinking at the ultimate level.

In the cultural case, transmission is achieved by human action and human thought. Humans are very far from indifferent about the contents of their acts and thoughts. The function of DNA replicase is to replicate DNA, but the function of humans is not to replicate culture. The function of humans is to be humans. This is a crucial difference. It means humans have all kinds of characteristic interests, strategies, goals, biases, priors, intuitions, and so forth. They apply shaping forces to whatever they transmit, sometimes unconscious and automatic ones, sometimes deliberate and reasoned ones; sometimes through individual action, sometimes through institutions. If it seems like humans unreflectively replicate just whatever society is doing, that's only because we focus on atypical cases. In the case of which side of the road to drive on, for example, my *only* interest in choosing one is to choose the same one as everyone else in the country. So for that trait, humans are remarkably obedient replicators of the society they live in, and a basically arbitrary norm is stable indefinitely. But most traits are not like that. Mostly, I have interests and biases that go beyond the mere desire for my behaviour to be the same as everyone else's[15].

The shaping forces in culture will be different for every trait. So too will the relationships of ancestor to descendant: you learn different things in different ways. So too will the dynamics: you learn somethings once and for all; others you continuously update through your life. There is thus no very general expectation we can form, for example that humans will hold such beliefs as maximise their cultural fitness, or anything like that. This is not a counsel of despair. I am not just saying 'it's all very complex', or that we are limited to a kind of post-hoc or qualitative historical interpretation of cultural change, without hope of bringing under the umbrella of natural science. Some of the transmission experiments described above show that this is not so. We can still formulate and test explanatory causal principles for the properties and dynamics of human cultures. But we need to begin from an appropriate framing of the problem in order to do this. I don't think that the failure for the 'beanbag' move to become widespread for culture in the way it did for population genetics is due to the perversity of social scientists. I think it is due to substantive and interesting differences between the genetic and cultural cases.

So am I saying, in short, that there can be no such thing as cultural evolutionary theory? There can certainly be cultural evolutionary *theories*. One can—indeed, one must—model cultural regularities and cultural change as the population-level emergent consequences arising from the ways individual people learn, communicate, influence one another, think, remember and forget. And formal, computational and empirical tools are required in order to do this. Much of the scholarly enterprise known as 'cultural evolution' is simply the attempt to provide these tools, without necessary commitment to the idea that cultural change is a narrowly Darwinian process. My worries here leave this work and its motivation intact. But note the difference from the genetic evolution case. Whereas

---

[14] See for example Fawcett, T.W., S. Hamblin and L.A. Giraldeau (2012). Exposing the behavioral gambit: the evolution of learning and decision rules. *Behavioral Ecology* 24: 2-11.

[15] For an approach to cultural evolution in whose spirit this paragraph was written, see Morin, O. (2015). *How Traditions Live and Die.* Oxford: Oxford University Press.

in genetic evolution, empirical knowledge of the transmission mechanisms in each case can be bypassed, in the cultural case, empirical knowledge of the transmission mechanisms in each case is *precisely what you need to have* in order to be able to build your explanation. You can't possibly understand how the repeated application of human cognitive biases shapes religious ideas without first studying in detail what those cognitive biases are. And your theory will only be as good as your characterisation of the cognitive or interactive processes that are doing the explanatory work. Theory is therefore dependent on knowledge of the transmission mechanisms in each individual cultural case, in a way that it is (arguably) not for genes.

It follows that there will be as many cultural evolutionary explanations as there are domains across which human cognition and human interaction are different. Thus, there can be *a* cultural evolutionary theory only in a weak sense, meaning the general set of recipes used in such explanations. These recipes will actually be rather varied, and will only work when made up with fresh ingredients from empirical psychology, cognitive science, politics, sociology, and so forth. This is rather different from the somewhat stronger sense in which there is *an* evolutionary theory in organismal biology. And that means if we wish to unify the social science disciplines, which surely we must wish to do, we must develop a slightly different banner under which to do it. I would put human action, or more generally human cognition, rather than cultural selection, at the heart of that enterprise.

I worry that I have spent a lot of this meditation being rather negative, so perhaps I should end instead by being positive: about humans, and about culture. Culture is not like DNA. It is the residue of past cognition and past interaction; a residue that is available to cognitively complex, socially interacting, purposeful, reasoning beings. And we are really good at doing things with this residue. The operations we perform on it are not limited to its reproduction. We can perpetuate it, yes, but also extend it, adapt it, discuss it, contest it, refuse it, restructure it, or redesign it; not in its interests, but in our own. The human capacity for purposive agency using the raw material embodied in culture is the resource a sensible progressive politics needs to be built upon. And it is the capacity that makes human culture, and indeed human beings, extraordinary.