Cultural Evolution Workshop

The International Cognition and Culture Institute

August-September 2010

On May 2010, Daniel Dennett gathered, at the Santa Fe Institute, a handful of people who have written about cultural evolution. The general impression was that (as he tweeted some time later) "the meeting revealed a lot of unexpected common ground". The International Cognition and Culture Institute published, by way of proceedings, each participant's

TABLE OF CONTENTS

PERSPECTIVES ON CULTURAL EVOLUTION

By Daniel C. Dennett

Ever since Darwin's Descent of Man (1871), the idea of adopting an evolutionary perspective on human culture has seemed to many to be a natural move, obviously worth trying—and to many others to be a dangerous, "nihilistic," "reductionistic", "scientistic," assault on everything we hold dear. Work on cultural evolution has been making good progress in recent years, but has been hindered by distortions, some perhaps deliberate, but others are misunderstandings that naturally arise between slightly different traditions. I formed this working party to try to find common ground and resolve differences among some of the leading theorists and experimentalists. The ten participants included the trio of Boyd, Henrich and Richerson (BRH), a French trio of Sperber, Claidière and Morin (SCM), the memeticists Blackmore and myself, and two philosophers of biology who have been particularly engaged with issues of cultural evolution, Peter Godfrey Smith and Kim Sterelny. Several other leading figures were invited but could not participate for various reasons.

Each participant was invited to send in two or three recent papers or chapters for everyone to read in advance — the list of these papers is available here –, and then the first three days were devoted to the "X on Y sessions", in which each participant (X) in turn took on the task of briefly introducing the work of another participant (Y). I invited all to send me their preferred list of people to introduce, and more or less optimized the pairings to make sure each X-Y pair were drawn from different traditions and no two introduced each other's work. After fifteen or twenty minutes introduction, each Y then had a chance to respond, followed by general discussion. The atmosphere was informal, permitting frequent interruptions for questions and comments.

Before the working group convened there was some skepticism and grumbling about the X on Y obligation from various participants, but everybody graciously acceded to my request and the results, in my opinion, confirmed the value of the practice. After the workshop all participants submitted a brief summary of the week, citing what was learned, what was agreed upon, and issues still unresolved. Quoting a few comments from participants: Peter Richerson: "I do think that the disagreements among the various 'schools' of cultural evolution represented at the meeting are relatively modest."

Peter Godfrey Smith: “I think that a lot of progress was made on clarifying disagreements, even where the remaining disagreements remain genuine […] It’s progress when an initially cloudy situation gives way to a sharper and more definite set of empiricial uncertainties.” Dan Sperber: “It has been a wonderful workshop of serious, demanding, insightful, informal, friendly discussion of a kind and quality rarely experienced.” Nicholas Claidière noted that part of the distortion is generated by the way we tend to talk about our work to people outside the field, giving the (wrong) impression that there are schools of thought at war with each other: “Given the amount of agreement that we have seen during this meeting, I think it would be more productive to present ourselves as having a common goal with diverging interests rather than competing views of the same phenomena.”

Terminological headaches

Three frustrating terminological problems were exposed, but we didn’t resolve how to correct them: “cultural group selection,” “meme,” and “Darwinian” are all good terms, historically justifiable and useful in context, but by now all are so burdened with legacies of ideological conflict that any use of them invites misbegotten “refutation” or dismissal. Should we abandon the terms in favor of emotionally inert replacements, or should we persist with them, always accompanying their use with a wreath of explanation? These are questions of diplomacy or pedagogical policy, not serious theoretical issues, but still, alas, unignorable.

As Boyd explained, the adoption by BRH of the term “cultural group selection” had its roots in the relatively uncontroversial theoretical terrain of Sewall Wright’s population genetics (and shifting balance theory), not in later, more dubious and controversial variants. But this is hard to explain to people who have already taken sides for or against “group selection” as an important phenomenon in evolution. In any event, the working group, enlightened about what BRH mean—and don’t mean—by cultural group selection, while still harboring somewhat different hunches about its importance, acknowledged that Steve Pinker’s recent “extreme and dismissive” (Henrich) position on Edge.org did not find a target in the work of BRH.

The popular hijacking of Dawkins’ term “meme” for any cultural item that “goes viral” on the Internet, regardless of whether it was intelligently designed or evolved by imitation and natural selection, has been seen by some to subvert the theoretical utility of the term altogether. There is also the unreasoned antipathy the term evokes in many quarters (reminiscent of the antipathy towards the term “sociobiology” that led to its abandonment). Alternatively, if one is “Darwinian about Darwinism” we should expect the existence of cultural items that are merely “memish” to one degree or another, and we might as well go on using the term “meme” to refer to any relatively well-individuated culturally transmitted item that can serve as a building block or trackable element of culture however it arrives on the scene. Other terms, such as Boyd and Richerson’s “cultural variant”, have been proposed, but the term “meme” has become so familiar in popular culture that whatever alternative is used will be immediately compared to, identified with, assimilated to meme(a Sperberian attractor, apparently), so perhaps the least arduous course is to adopt the term, leaving open its theoretical definition, in much

the way the term “gene” has lost its strict definition as protein-recipe in many quarters. Since the long-term fate of such an item will be settled by differential reproduction (or something similar to differential reproduction) however much insight or “improvisational intelligence” went into its birth, it has a kind of Darwinian fitness.

But should we go on talking about whether or not a phenomenon is “Darwinian”? Some think the term gets in the way, since we are seldom if ever alluding to what Darwin himself thought, but rather to the neo-Darwinian, post-DNA synthesis, itself an evolving landmark. On the other hand, there is general agreement within the group that some important elements of human culture evolve by processes strongly analogous to genetic natural selection, and the variations in these processes can be usefully diagrammed using Peter Godfrey Smith’s “Darwinian spaces” (See figure 1 for an instance), in which the similarities and differences can be arrayed in three dimensions. Since, moreover, there is agreement that these cultural regularities can set selection pressures (e.g., a “cultural niche”) for co-evolutionary processes, generating genetic responses (such as adult lactose tolerance), a unified evolutionary perspective, in which the trade-offs between cultural and genetic evolution can be plotted, is a valuable organizer of phenomena, some “more Darwinian” than others. No other term suggests itself for the set of features that mark paradigmatic (neo-)Darwinian phenomena, so perhaps the misunderstandings the term tends to generate can be deflected.

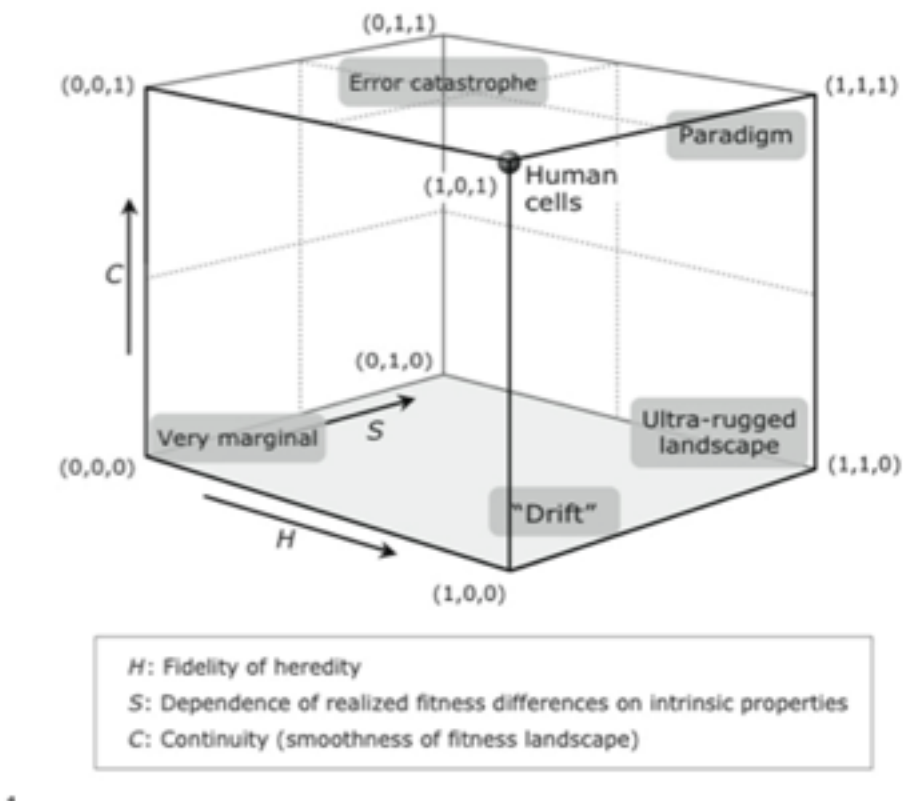

Figure 1

Consensus

The working group agreed on a number of points, some methodological and some substantial, that are still considered controversial by others, or in some cases just not yet considered:

1. We should be Darwinian about Darwinism; there are few if any bright lines between phenomena of cultural change for which cultural natural selection is clearly at work and phenomena of cultural change that are not at all Darwinian. The intermediate and mixed cases need not be marginal or degenerate, a fact nicely portrayed in Godfrey-Smith's Darwinian Spaces.

2. Models must always "over-"simplify, and the existence of complications and even "counterexamples" relative to any model does not automatically show that the model isn't valid when used with discretion. For instance, the absence of explicit treatment of SCM's "hetero-impacts" in BRH's models "does not amount to a denial of its importance"(Godfrey-Smith). Grain level of modeling and explaining can vary appropriately depending on the questions being addressed.

3. The traditional idea that human culture advances primarily by "improvisational intelligence," the contributions of insightful, intentional, comprehending individual minds, is largely mistaken. Just as plants and animals can be the beneficiaries of brilliant design enhancements that they cannot, and need not, understand, so we human beings enjoy culturally evolved competences that far outstrip our individual comprehension. Not only do we not need to "re-invent the wheel," we do not need to appreciate or understand the design of many human institutions, technologies, and customs that nevertheless contribute to our welfare in various ways. Moreover (a point of agreement between Sperber and Boyd, for instance), the opacity of some cultural memes (their inscrutability to human comprehension) is often an enhancement to their fitness: "This opacity—which is a matter of degree, of course—is what makes social transmission so important. It plays, I believe, a crucial role in the acceptability of cultural traits: it is, in important ways easier to trust what you don't fully understand and hence cannot properly evaluate on its own merits." (Sperber)

4. The persistence of cultural features that are not fitness-enhancing, and may even be fitness-reducing, is to be expected in cultural evolution, and can have a variety of explanations.

New questions

1. Rob Boyd, in his post-working group summary, proposed a way in which the Evolutionary Causal Matrix idea developed by Sperber and Claidière can be re-expressed in the population genetics formalism used by BRH, raising questions about how—if at all—the homo-impact/hetero-impact distinction introduced by SCM appears in the population genetics formalism. Do SCM have a reply? [1]

2. SCM propose that cultural attraction, not differential replication, accounts for much of the dynamics of cultural evolution [2] (in the neutral sense: change over time), but several expressed concern that only a (quasi-)Darwinian process can initiate and refine adaptations (lifting in Design Space). One line of thought suggests that attraction and replication can sometimes work together: attractors act rather like norms to somewhat digitize otherwise continuous variations, making exemplars stable and distinct enough to be eligible for iterated replication and selection. Another line of thought is that the distinction between attraction and differential replication is maybe just a question of "zoom": if you

zoom in on apparent replicators, you may find that they are not, strictly speaking, replicating at all, but if you zoom out, the results are as if there was replication going on. [3] Which of these suggestions will survive further research? For instance, are there experiments (Claidière's question) that can distinguish the roles of transformative and selective processes, shedding light on the conditions under which each plays the dominant role?

3. "If individuals are smart enough in their choices, the BRH meso-level picture fades. When people are smart and make good choices, the recurrence of good options and accumulation of design can occur without imitation-and-selection." (Godfrey-Smith) But Sperber points out that this need not pose a dichotomous choice between evolutionary and rational-choice explanations: "adding attraction to the cultural evolution story allows us to integrate evolved mechanisms that tend to produce rational choices, not as an alternative kind of explanation, but as a factor of attraction among many." Under what conditions can this proposed unification do serious explanatory work? Since attractors can be both enhancers and decelerators of adaptive change, are they too versatile to be explanatory (at least in this context)? [4]

4. Is cultural evolution "de-Darwinizing" (Godfrey-Smith's term for phenomena that evolve into less Darwinian phenomena)? Dennett says yes: in the earliest days of human cultural evolution, individuals were largely uncomprehending beneficiaries of their new tools and customs, only gradually becoming reflective, critical, foresighted users of those tools. Today they aspire to be intelligent (re-)designers of every aspect of their environments, and some of the major changes in culture today are the products of quite concentrated, not distributed, R&D. [5] Blackmore says no: on the contrary, technology has raised the proportion of high-fidelity copying and transmission, and is beginning to usurp the role of the supposedly intelligent designer thanks to automated search and evaluation systems. Will all roles for human "improvisational intelligence" become obsolete, and "inventors" as rare are telephone operators, coopers, and scythe-sharpeners in the future? Or will the heretofore unreachable ideal of the intelligent designer be approximated by individual human beings, thanks to their reliance on technology (including especially instruction and the cascade of scientific knowledge that creates new platforms from which to begin one's exploration)? Human civilization today appears to be a volatile mix of these opposing trends; are there investigations that can clarify the resultant direction in which we are heading?

5. Richerson raises an issue (among many others) that we did not have time to discuss: "Natural selection on genes admits of a number of modes. […] Throw in density and frequency-dependent selection. […] Mate choice and artificial selection introduce agent-based rather than natural selection, demi-god designers if you want. With cultural evolution agent-based social selection runs wild." Does this point to a good way to organize the intermediate space between paradigmatic "Darwinian" natural selection and intelligent design? One thing that is changing in this progression might be called the focus of the selection pressure. At the Darwinian pole (simple natural selection) the selection pressure is "just" a statistical net effect of a kazillion independent events that determine which candidates

get replicated; in the middle-ground, mate choice (as Geoffrey Miller has argued) is focused through the perceptual/cognitive/emotional dispositions of individual (usually female) "minds," with varying degrees of comprehension and reflection; it is like Darwin's "unconscious" selection which bridges the gap between agentless natural selection and reflective, intentional "methodical" selection. As agents (conceived as mere concentrations of selective efficacy, selective "hot spots" in the environment) [6] become more discerning, the importance of high-fidelity replication does not lapse, but the breadth of "search" contracts and R&D can become more efficient (it can also hasten the ruin of ill-informed R&D). As reflectivity about this very process increases, R&D becomes faster and more efficient—but gradually, allowing for opaque attractors to play a large role relative to genuinely insightful or comprehending quality judgments. Does this proposal withstand scrutiny?

Comments by Peter Richerson and Dan Sperber

Richerson commented on the draft of this document and Sperber replied:

[1] Richerson: I thought that the attraction concept had become sufficiently generalized as to obviate this distinction. Perhaps complete resolution of this issue need to await SCM's development of their models. With a fully functional model in hand, we can see if the structure of them differs in some fundamental way from the population genetics based models I'm more familiar with.

Sperber: My first reaction to Rob's comments was, to begin with, sheer joy at having him discuss ECM seriously. Given Rob's experience and competence, this cannot but be good for the science. Were Rob to find that there is a basic flow in the ECM approach, then we would be spared going in the wrong direction, and again, good for science. Rob might also find ways to correct and improve the ECM format at least for some use, and this would be nice, of course.

Now, regarding, the fact that "the ECM formalism can be equivalent ways of representing exactly the same underlying processes," I like Rob's illustration, and Nicolas and I had found other examples in our work in the past. I don't see this as an objection, especially since we didn't propose the ECM format as it stands as an alternative way to model population phenomena of interest, let alone as a better way. We offered it as to begin with a Dennettian 'intuition pump', leaving open the question whether it could, at least in some cases, be developed into a perspicuous way of modeling. The intuition pump effectiveness was, for me, demonstrated at our workshop and in several other exchanges I have had: people who didn't quite 'get' the attraction idea, found it much easier and even congenial when so presented.

On the further more technical points raised by Rob, I would like to coordinate at least with Nicolas and Thom before providing a careful reaction.

[2] Sperber: What we propose is that hetero-attraction is likely to be more or much more than a marginal factor in cultural evolution, making a generalized notion of attraction that includes both homo and hetero-attraction – I agree with Pete with his comment on this point – potentially quite useful. This by

itself does not determine which is the best way to model cultural evolution, or precludes the possibility that different models may be better for different types of cases.

[3] Sperber: Here I agree with a remark Rob made in his comments: yes we, the attraction people tend to zoom towards greater details, but this doesn't necessarily preclude the possibility that on some issues at least, a more standard population genetics provides for a better zoom.

[4] Sperber: Here you want to talk about specific factors of attraction and the way they may contribute to adaptiveness, or to the resilience of non-adaptive features. The relevant point here is that the evolved ability to recognize and, under certain conditions, even design well-adapted things is a powerful factor of attraction that contribute to explaining the cultural success of well-adapted things. You get your evolutionary explanation, as usual by looking at micro-processes at a population scale. The fact that, in this case, rational choice modeling can also make the right prediction does not in any way undermine a more standard evolutionary approach (that moreover does better at least in terms of generality and of psychological plausibility).

[5] Richerson: Nuts Dan! Highly innovative places like Silicon Valley are Darwinian pressure cookers. First, the finest engineering training available in the world dumps the max amount of accumulated wisdom into the heads of the best and brightest. Then the B&B are set to work finding marginal impro-vements in existing designs to patent. Entrepreneurial teams funded by venture capitalists recombine old designs and add the latest new patented ideas to create products that are selected in ruthlessly competitive markets.

Dennett responds: But this Darwinian "pressure cooker" is distant from the Darwinian paradigm in several important dimensions: it is what Darwin himself called "methodical selection" (in his wonderful introductory passage that segues from the (intelligent) selective actions of plant and animal breeders, through the "unconscious selection" of the inadvertent, or largely purposeless biases of human beings in the early days of agriculture, to "natural selection" (in which no mind, intelligent or clueless, is required). The search space is pinched by many preconceptions, good and bad, and, as in sexual selection, the winners have been aggressively tested by nervous systems tuned to detect quality.

[6] Sperber: Yes, let's not overdo 'agents'. '"Hot spots" in the environment' is a nice metaphor. Another, more detailed way to go is to see cognition both as massively modular and heavily situated/distributed. At this point, the individual organism is still in play, but most cultural phenomena are both infra- and trans-individual (or to use Dennettian terms, sub-personal and collective) The agents that rational choice theorist theorize about not only don't exist – that is not too bad –, they are not, I believe, a very good idealization for modeling cultural evolution (this might be a point of difference between the attraction approach and the agents-choose-variant approach).

THOUGHTS ON THE WORKSHOP (1)

By Susan Blackmore

There was much agreement at the wonderful working party in Santa Fe. For example, we agreed about the importance of re-production, reconstruction, teaching and demonstration as well as true imitation. I loved DS's T-shirt folding video, but concluded that the variety and complexity of these processes does not detract from the fact that the folded T-shirt and the skill of folding it are memes that are passed on and selfishly compete with other variants.

We all agreed that there is always a leash (if have previously implied that the memes could entirely escape, I would now suggest this is only possible in the digital world of temes). We all agree that, as RB put it 'imitation did not just happen' but arose because of relevant selection pressures.

I'm sure others will describe the many areas of agreement better than I can. I would prefer to note some points of disagreement or openness and the shifts that took place as we discussed the wide range of differing ideas and research methods.

In my introduction I wrote, "The question that drives my interest is whether memetics is fundamentally different from other theories of cultural evolution or not." By the end of the week I am much clearer about the issues involved but concerned that to some extent we need to resolve semantics first. For example, some (e.g. OM) would like the word 'meme' to expand to include all cultural items (which would include some that recur without any kind of replication), while others (NC) would like to stick to memes as replicators. I remain convinced of the power of memetics while others are not. For example, RB and DS argue that it should be seen as a special case of wider cultural processes and PR emphasises the many items that are not at all meme-like or 'memish'. Here are some further thoughts.

Be careful of analogies

NC and DS claim that for memetics the analogy between genes and memes 'is deep indeed' and assumes direct equivalents, such as a cultural phenotype. I argued that memetics is not based on analogy but on the principle of universal Darwinism: the idea that memes undergo the same evolutionary

algorithm as genes. This means some analogies will be close (because of the underlying processes) and others misleading (because genes depend on high fidelity cellular chemistry while memes depend on the complexities and weaknesses of human meme machines). The most troublesome analogies are:

1. Germ-line phenotype distinction. RB suggested that memes live in vehicles or interactors. But for most of memes' relatively brief life there has been no germ-line phenotype distinction and so no meme vehicles or interactors. However, as one might expect, they have recently appeared and are spreading fast. Printing presses, car factories and computer software all copy the instructions for making more books, cars and digital products rather than copying the products directly. Kayaks (an example much used at the meeting) were copied from other kayaks during most of their evolution but when I bought my modern plastic kayak in a shop, I tried out several models before choosing one of many identical ones made on a factory production line.

2. Guided variation or directed change. This disanalogy is often remarked upon, especially in criticisms of memetics. 'Learners are smart shoppers not compulsive imitators' (PR on OM). So it was useful to hear about instances of possible directed change in biological evolution. I was previously under the impression (from many comments in the literature) that guided variation is always destructive but having learned more from our discusssions I now want to think about when and how it either speeds up the search or leads to traps. I would expect the digital revolution to provide examples of increasingly random variation but I do not know of any non-trivial examples of this.

Are memes attractors?

DS urged me to agree that memes are attractors. I much enjoyed reading the Phil Trans paper, presenting 'Blackmore on Claidière', and using ECMs to think about cultural evolution, but I became increasingly concerned about where innovation comes from. A breakthrough occurred for me when DS explained that the matrix applies to memes that do not yet exist, leading DD to imagine new memes being 'sucked into existence' by hetero-impacts and a discussion of whether this entails cranes or skyhooks, and how shifting the level of abstraction can switch homo- to hetero-impact. In the end I conclude that ECMs describe an abstract space in which memes are drawn into attractors, but memes themselves are not attractors.

Are memes information?

Yes. PGS gave us an excellent primer on information and its bearing on sending, receiving and copying information, and made the interesting point that true imitation requires no co-evolved sender-receiver relationship. After interesting discussions I concluded (for the moment) that, as DD put it, 'all the push and pull happens at the level of tokens' but for explanatory purposes we need to work at the type level (e.g. discussing words as memes). Not everyone agreed about the status of either genes or memes as information but nearly everyone used the word 'information' when discussing memes. I

think it's still helpful to stick with Dawkins' original formulation of memes as information copied with variation and selection, even though so much uncertainty surrounds what is meant by 'copied' in this context.

Are memes replicators?

This is really the crux for memetics. I may have been guilty of emphasising too strongly the role of imitation. So now, having been urged to think a lot more about the relevance of reconstruction, re-production, teaching and demonstration, I will be more careful. Yet I agree with OM that in thinking about memetic evolution the psychological complexities can, for some purposes, be abstracted away (or bracketed), and cultural items that are propagated by many different methods can still be classed as replicators. This still leaves the question of how much of culture can be considered meme-like. NC's animal work was useful here, as were the discussions of cultural group selection (RB, PR). On this topic there was a lot of discussion and little agreement – not, I think, because people had differing fixed views but because of the complexity of the issues. I think we made a lot of progress even if we came to no firm conclusions.

Creativity

DD claimed that the termite mound and the cathedral, though looking similar, are created by entirely different processes. I believe they are really the same. The cathedral depends upon meme evolution both between people and within the architects' and builders' heads. So both constructions are designed by Darwinian processes (is there any other kind of design?). This relates to OM saying there is no core; no inner self as designer. Theories of cultural evolution extend the tendency of science to overthrow our self-centred view of the universe.

Where are memes going?

DD claims that culture is getting less Darwinian, whereas I think the opposite. Many (perhaps all) of us found PGS's 3d spaces helpful, and he and DD gave interesting examples of more or less Darwinian cultural processes. However, I cannot see that, in general, culture is de-Darwinizing. Recent increases in the fidelity, longevity and fecundity of memes suggest the opposite, e.g. with digitisation through language, writing, printing and computers, and the cultural shift from learning by apprenticeship to learning by text. As for the future, I predict further increases and greater autonomy for memes, as well as a shift in power towards a third replicator based on digital information copied, varied and selected by machines. Such temes will surely continue spreading and increase in their ability to restructure the human mind.

I found the whole week absolutely delightful – challenging, interesting, and exciting. I would like to thank everyone who was there, as well as Dan for bringing us all together in such a wonderful place.

THOUGHTS ON THE WORKSHOP (2)

By Rob Boyd

I have been thinking about the relationship between the kind of models that Pete, Joe and I have made and the ECM framework sketched by Dan and Nicolas in their paper. In particular, I have trying to understand how the mechanistic processes represented in our approach relate to homo-hetero distinction central to the ECM approach. In my remarks on Sue Blackmore's papers I sketched a simple model in which there were two variants with a selection stage and a transformation stage. So, first I'd like to convert this model to the ECM framework developed by Sperber and Claidière. Label the variants 1 and 2. Let W_i be the fitness of variant i in the selective stage. Individuals who acquire variant 2 remain variant 2, but individuals who acquire variant 1 transform to variant 2 with probability m: Like Nicolas and Dan, I will assume that populations are large enough that drift-like sampling processes can be ignored. These assumptions lead to the following recursions for the frequencies of the two variants, f_1 and f_2 :

$$f'_1 = \frac{x}{y}\frac{(1-m)f_1W_1}{f_1W_1 + f_2W_2}$$

and

$$f'_2 = \frac{mf_1W_1 + f_2W_2}{f_1W_1 + f_2W_2}$$

Since $(1-m)f_1W_1 + mf_1W_1 + f_2W_2 = f_1W_1 + f_2W_2$ these recursions are equivalent to the following Evolutionary Causal Matrix

$$\begin{pmatrix} (1-m)W_1 & 0 \\ mW_1 & W_2 \end{pmatrix}$$

This exercise illustrates that the population genetics/epidemiology formalism and the ECM formalism can be equivalent ways of representing exactly the same underlying processes, and that the transformation of variant 1 into variant 2 appears as a hetero effect of variant 1 on the frequency of variant 2. Notice that variant 2 has no hetero effect on variant 1. This suggests, to me at least, that the magnitude of the coefficients in the ECM may not represent the causal processes involved in cultural evolution. For example, it could be that the variants are two versions of a story, and learners have existing theories that transform a fraction m of all stories into variant 2. A fraction of the learners who hear variant 1, think they have heard variant 2. This seems to me best thought of as variant 2 having a hetero effect on 1, not the reverse as the coefficients suggest. On the other hand, it might be that variant 1 is more complicated and harder to remember, so that learners learn 1 but sometime later remember 2. This seems to me to be best thought of causally as an effect of 1 on 2. I'd be interested to know what Nicolas and Dan think about this.

I also think it's important to see that even very simple models will require ECM's in which matrix elements to depend on variant frequencies. Consider a second, even simpler model which only contains transmission and transformation (a guided variation model in the RHB jargon and analyzed in Boyd and Richerson 1988) no selection at all. Suppose individuals have an individual learning mechanism that responds to environmental cues. There are two variants again labeled 1 and 2. With probability p_i they receive a cue that tells them that variant i is best and they adopt that variant. With probability $L = 1 - p_1 - p_2$ the cue does not clearly tell them which variant is best and they imitate a randomly chosen individual from the previous generation. This leads to the following recursions for the frequency of the two variants:

$$f'_1 = Lf_1 + p_1$$

$$f'_2 = Lf_2 + p_2$$

First notice that the population evolves to an equilibrium at which the variant that is more likely to be learned is at higher frequency $\hat{f}_1 = \dfrac{p_1}{p_1 + p_2}$ and thus if $p_1 \gg p_2$ transformation alone can lead to cumulative adaptation. Also, notice that since there is no selection in this model, there is no need to normalizefrequencies by dividing by an average fitness like term. Thus the ECM is

$$\begin{pmatrix} L + \frac{p_1}{f_1} & 0 \\ 0 & L + \frac{p_2}{f_2} \end{pmatrix}$$

Since the off-diagonal elements are both zero, there are no hetero effects in the model at all. The L components of the homo terms are fine; they just represent incomplete transmission. But the second part of these terms are peculiar in that they are inversely proportional to the frequency of the relevant variant. This is formally necessary in order to represent frequency independent learning, but it doesn't seem to reflect the causal structure of the processes modeled. What is really going on is that there is a frequency independent learning process that creates variants one and two with probabilities $p1$ and $p2$. This suggests that the ECM frame work should be extended by adding a frequency independent term that represents the effect of the non-cultural environment on the frequencies of the cultural variants.

Finally, I think that it will be important to think carefully how to incorporate multidimensional cultural variants into the ECM framework. Consider a trait that has two dimensions. Each dimension can have two states which I will imaginatively label 1 and 2. For example, dimension 1 could the the length of the bow (long or short) and dimension 2 could be ether it is sinew backed (yes or no). So there are four variants 11, 12, 21, and 22. Individuals squire both dimensions form one of their parents. This cultural trait has no effect on the probability of becoming a parent and there is no error, so that transmission leaves the frequency of the four variants unchanged, and thus the ECM matrix is

$$\begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}$$

The influences are purely homothere is no interaction between traits and no hetero influences.

Now let's modify the model a bit. Suppose that with probability r individuals individuals acquire the two dimensions of their behavior from two different randomly chosen individuals, and with probability $1 - r$ they acquire both dimensions from the same individual. Thus, for example, the frequency of variant 11 in the next time period is

$$\begin{aligned} f'_{11} &= r(f_{11} + f_{12})(f_{11} + f_{21}) + (1 - r)f_{11} \\ &= r(f_{11}^2 + f_{11}f_{21} + f_{12}f_{11} + f_{12}f_{21}) + (1 - r)f_{11} \\ &= r(f_{11}(1 - f_{22}) + f_{12}f_{21}) + (1 - r)f_{11} \\ &= f_{11} - r(f_{11}f_{22} - f_{12}f_{21}) \\ &= f_{11} - rD \end{aligned}$$

The population genetics student will recognize this as a recursion giving the effect of recombination on gamete frequencies. The recursions for the other three cultural variants are

$$
\begin{aligned}
f'_{12} &= f_{12} + r(f_{11}f_{22} - f_{12}f_{21}) = f_{12} + rD \\
f'_{21} &= f_{21} + r(f_{11}f_{22} - f_{12}f_{21}) = f_{21} + rD \\
f'_{22} &= f_{22} - r(f_{11}f_{22} - f_{12}f_{21}) = f_{22} - rD
\end{aligned}
$$

This set of recursions can be represented by many distinctive ECM's. To see this, let's concentrate on the first row of the ECM which is

$$1 + rf_{22} \quad -rf_{21} \quad 0 \quad 0$$

or

$$1 + rf_{22} \quad 0 \quad -rf_{12} \quad 0$$

or

$$1 \quad -rf_{21} \quad 0 \quad rf_{11}$$

or

$$1 + rf_{22} \quad -\tfrac{1}{2}rf_{21} \quad -\tfrac{1}{2} \quad f_{12}0$$

And there an infinite number of possibilities. The problem is that the terms proportional to products of frequencies like $f_{12}f_{21}$ can be represented in two different columns of the matrix, or as any convex combination of the two terms (i.e. the weights have to sum to one). I think this means that it is not possible to read the hetero/homo effects directly from the matrix, but I am not sure.

Perhaps some convention could be established that would allow the matrix to accurately reflect the causal process of partial recombination.

I am also puzzled by another aspect of this model. Let $D = f_{11}f_{22} - f_{12}f_{21}$. This is the covariance between the trait value along dimension 1 and dimension 1. For example, if short bows are more likely to be sinew backed $D > 0$. Let p_i be the frequency of variant i (long or short) for dimension 1 and q_j be the frequency of variant j (backed or not) for dimension 2. We can write down recursions for the frequencies of each dimension, and for D.

$$
\begin{aligned}
p'_1 &= p_1 \\
p'_2 &= p_2
\end{aligned}
$$

which is equivalent to the ECM

$$\begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$$

which represents pure homo effects. The ECM for the q_i's is the same and there is a recursion for D, $D' = (1 - r)D$. It is not easy for me to see how to incorporate the later recursion into the ECM framework. If we added selection-like processes to the model, I believe that they would end up as homo processes in the first representation (one trait, four variants) and as hetero processes in the second representation (two traits with two variants and a covariance) in which the degree of hetero influence would be proportional to the covariance divided by a variance (i.e. how much a variant on of one trait predicted the variant of the other trait).

At the end of the meeting I commented to Dan that Pete, Joe, and I tended to look at these processes at coarser scales than he, Nicolas, and Olivier. And, I still think this is true in some ways. The latter three are more interested in the micro causal details that give rise to persistence. For example, Pete, Joe and I have thought of there being alternating cycles of internal and external representations. However, we preferred to zoom out and try to write models that black boxed the details of this processes while Dan, Nicolas and Olivier want to think about the causal processes in detail. However, the current exercise makes me think that there are also ways in which the kinds of models that Pete, Joe, and I have made are less coarse grained than the ECM approach because they build in more meso scale detail about the dynamic processes.

Reference

R. Boyd and P.J. Richerson (1988). An Evolutionary Model of Social Learning: The Effects of Spatial and Temporal Variation. In: *Social Learning: A Psychological and Biological Approaches*, T. Zentall and B. G. Galef, eds., Lawrence Erlbaum Assoc., Hillsdale, NJ.

THOUGHTS ON THE WORKSHOP (3)

By Nicolas Claidière

(Clarification: I use disagreement in the sense of interesting topic that could be discussed/ researched further and questions in the sense of more precise ideas that could help sort out disagreements.)

Interrogation: It just dawned on me that we could have discussed a, maybe important, academic matter: that of our presentation to other academics. When talking about our respective work we often refer to alternative theories (e.g. dual-inheritance, attraction, memetic, etc.) which I think gives the wrong impression. Given the amount of agreement that we have seen during this meeting I think it would be more productive to present ourselves as having a common goal with diverging interests rather than competing views on the same phenomena. A first simple way of achieving this is would be to not present our respective work in terms of alternative theories but just refer to actual articles (so instead of saying 'memetic theory assumes that', I am now going to say 'Dan Dennett told me that'). Another, more complicated way would be to agree on a common denomination (I am really bad at the naming business so I won't even try to make a suggestion here) that could for instance figure in the title of the report Dan will prepare and that we might want to publish (I personally would like that). Anyway, I thought I would throw this out because I would like to hear what you think about that.

1. The populational approach to the study of cultural evolution

Agreement: I think we all agree that approaching cultural evolution as a population of cultural variants (aka meme, items, tokens, etc.) that are transmitted by individuals is insightful. I think we also all agree that the right ontology is at the token level for that kind of description and that zooming in/out of the token by token description provides different insights.

Disagreement: The extent to which the Populational view of cultural evolution is Darwinian is unclear. It could to some interesting extent not be Darwinian at all, in the minimal sense proposed by PGS. Or it could fit the minimal definition but be closer to marginal rather than paradigm cases (that's what I think). Or maybe this is not an interesting question anyway (RB).

Questions: An outstanding question with respect to the Darwinianism (or Darwiniality????) of cultural evolution is the role of multiple parents in generating offspring, thereby blurring genealogical relationships. One example is the case of language phylogenies and the possibility that they are an example of phylogenies without underlying Darwinian populations. Another project concerns the development of a PGS space for cultural evolution.

2. The origin of complex design

Agreement: Culture is adapted in the sense that some serious lifting in design space has been achieved and to some interesting extent this comes from cumulative cultural evolution, the gradual accumulation of cultural modifications over time. Social learning is essential to go beyond individual learning and explain the lifting but explaining exactly how social learning improves the lifting is still debated.

Disagreement: The origin of complex cultural adaptations can come from high fidelity copying (arising from various processes such as imitation, teaching and demonstration) associated to selective processes (view of RB, PJR, JH, DD, SB?) and/or it can arise through a combination of constructive processes, such as individual learning, associated with low-fidelity transmission (as in the apprentice model; view of DS, KS, OM, NC?).

Question: I think the burden of proof lies with the second view here, since it is clear that high-fidelity copying plus selection leads to adaptation. Experiments addressing the role of individual learning plus low-Fi transmission would be useful.

3. The balance between selection and transformation

Agreement: The evolutionary change between two time steps can be partitioned into transformative (aka directed variation, directed change, constructive processes, etc) and selective processes (through differential multiplication brought about by conformity, prestige, etc). The Price equation and RB's model of selection/mutation balance could provide a useful illustration here. The outcome of the evolutionary process depends on the relative strength of transformative vs. selective processes and is likely to vary substantially between cases.

Questions: To me experiments in which we can partition the evolutionary change between transformative vs. selective processes can provide valuable insight into cultural evolution.

4. The future of memetic

Agreement: The meme's eye view provides a useful perspective on cultural evolution and asking "who benefits?" can force us to take into account the fact that cultural evolution need not happen because it benefits individuals in any way. This is especially valuable when there is a conflict between the meme and the host.

Disagreement: It is not clear if memes are, or ought to be, replicators. Replication happens in cultural transmission but given the usual amount of directed changes that occurs, replication probably is a more marginal than paradigm case.

Questions: It is unclear whether culture evolves to become more 'replication like' or not.

5. Topics on which I am agnostic

Gene-culture coevolution: I think we have undeniable evidence of gene-culture coevolution (e.g. lactose, etc) but it is unclear to me to what extent gene-culture interactions matter in cultural evolution in general. For instance, I can imagine gene-culture coevolution being very important in cases related to disease resistance and food consumption, both having strong effects on biological fitness. I am more skeptical about other domains because I don't think that the cultural variation that has strong fitness effects will in general last long enough for biological evolution to act substantially. In other words, I think that for culture and genes to coevolve a certain number of special conditions have to be met and that this is not generally the case. That however, is more a hunch than a claim and I am not competent to discuss the archaeological/genetic evidence here.

Cultural group selection: Again, it seems to me that as presented by RB and PJR, there is no reason not to expect some cultural group selection but the amount and the role in cultural evolution is still unclear to me.

THOUGHTS ON THE WORKSHOP (4)

By Peter Godfrey-Smith

I think that a lot of progress was made on clarifying disagreements, even where the disagreements themselves remain genuine. Many of the remaining disagreements are empirical. It's progress when an initially cloudy situation to gives way to a sharper and more definite set of empirical uncertainties.

Micro, Meso, Macro. To set things up I'll make explicit some distinctions between levels of description – between coarser and finer grained perspectives on a cultural system.
Micro-level: I take this to involve individual psychology, person-to-person social interaction, and the making of artifacts by individuals.
Meso-level: Coarser-grained facts about a single culture or population. The spread of a new bow design or a new taboo would be examples.
Macro-level: Cultural phylogenesis and related events. A whole culture might split into two or go extinct.

Explaining adaptation and design at the meso-level. I take Boyd, Richerson, and Henrich to make their central claims at the meso-level. They argue that given the low levels of comprehension in individuals, with respect to the complex tools they use and their knowledge of the world, selectionist processes (in a very broad sense) must be important in culture. Populations are smarter than their consituent individuals. The increase in "R&D" (Dennett) in cultures is not a simple accumulation of intelligent moves by individuals. Instead there must be a general pattern of accurate transmission of variants, with good options being passed on more than others. Then there can be the accumulation of improvements by small steps. Dennett I take to agree with B, R, & H here. Perhaps Blackmore does too, though she would describe these changes in a different way. Sterelny is also sympathetic, but not so convinced that the level of comprehension is generally low.

Looking up from the micro-level. Sperber, in contrast, starts from the micro-level, and things look different from there. I think that Claidiere and Morin are on roughly the same page as Sperber here, so this package will be 'SCM.' SCM think that at the micro-level, there is less of a role for faithful copying

and a substantial role for psychological 'attractors.' People reconstruct cultural variants rather than copying them, and do so in a way guided by pre-existing psychological structures. As became clear in discussion, attractors need not derive from general features of human psychology. A previous round of cultural change can give rise to attractor-like constraints on how people reconstruct cultural variants in the next round.

Here is a question that became clearer without being resolved: does the SCM view of the micro-level conflict with BRH claims about the meso-level, or are they compatible?

Some micro-level stories that look antithetical to a Darwinian view of culture need not really be so. This is my interpretation of some of the S&C material about "hetero-impact." Cycles of hetero-impact look like something different from replication, but in fact the recurrence of cultural variants might be of the right kind to allow the BRH story to stand at the meso-level. As Boyd emphasized many times, when constructing a model, you must simplify and idealize in some places in order to reveal structure elsewhere. So the absence of explicit treatment of hetero-impact in the BRH models does not amount to a denial of its importance. A variety of possibilities at the micro-level are compatible with the central BRH claims at the meso-level. (See also Rob's summary, especially his final para, here.)

However, there are some micro-level possibilities that would make the BRH meso-level view less plausible. Here the role of comprehension is important. If individuals are smart enough in their choices, the BRH meso-level picture fades.

When people are smart and make good choices, the recurrence of good options and accumulation of design can occur without imitation-and-selection. Sterelny argues that recurrence of behaviors across generations at the micro-level often involves teaching and apprenticeship, and this is not a low-comprehension matter. So I took some of the debate on day 1 between Boyd and Sterelny about comprehension to be very important. This is an empirical debate. Morin's summary comments, and his talk, also make a plea for "improvisational intelligence" in at least some domains.

Memes. In this area I think there was a "move towards the middle ground" over the course of the week. Here is my attempt to make the middle ground explicit, combining elements from Dennett, Blackmore, and Sperber and linking them to the ideas above: meme-talk is appropriate as a way of discussing recurring cultural objects that are produced and used in a somewhat low-comprehension way. They need not be replicators, even in a relaxed sense. They might arise by hetero-impact. But there is (or should be) a real difference between a meme-based view of cultural variants and a traditional rational choice framework. It is not the case that memes are just whatever recurs in culture; Sperber pointed out that this would trivialize the meme framework, and I think that is right. But memes need not be copied. There might be a role for attractors, hetero-impact, and so on. So the viability of this relaxed view of memes is tied to the empirical disagreement described above about comprehension.

I guess I think that the future of meme-talk will be in informal summaries of low-comprehension processes of cultural change, rather than actual theory-building. To say that is to disagree with Blackmore and Dennett on the theoretical importance of the cui bono issue.

Spatial representations. Modifying Dan's initial chart, in my space I had three axes: comprehension, centralization, and saltation. Culture is more Darwinian when it has low values of all three. (And though there are plenty of populational and evolutionary processes that are not Darwinian, I think the chart captures something about the viability of evolutionary views of culture in a broader sense, too.)

I take it that the role of saltation is clear and uncontroversial, though there will be empirical disagreements. There was a lot of debate about comprehension, as I noted above. The locus of debate is especially the lower left "Kim versus Rob" part of the figure.

We did not discuss centralization very much, but I think everyone agrees that highly centralized societies are less amenable to an "evolutionary" treatment, in any non-trivial sense of that term. Something we did not discuss is the fact that a highly centralized society can still change by trial-and-error. In that case a different organic metaphor for culture becomes applicable, the idea of a culture as akin to an individual learning agent. Here's a slight update of my chart from the last day:

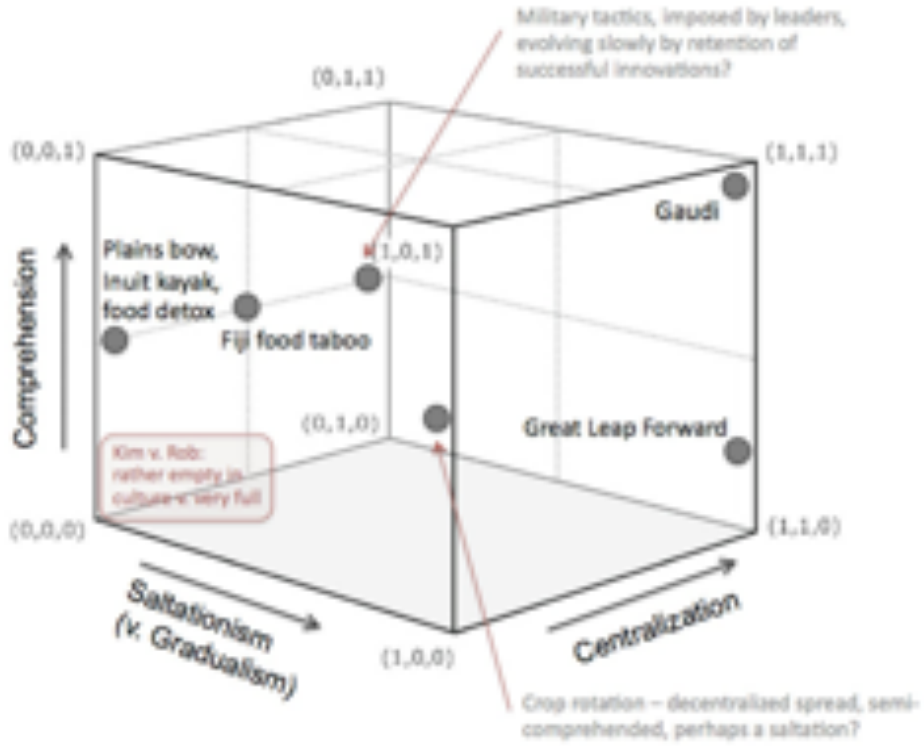

Cultural change is more Darwinian when it has low levels of Comprehension, Saltation, and Centralization

THOUGHTS ON THE WORKSHOP (5)

By Joseph Henrich

Based on our five days of discussion, I tried to come up with a list of stuff we agree on:

1. Cultural evolution is a crucial phenomenon for understanding humans (at least). It's a worthy goal to develop a broad framework for thinking and studying cultural evolution.

2. Natural selection has shaped human minds in ways that have a big impact on cultural evolution.

3. Some of these reliably developing psychological products of natural selection can usefully be thought of as adaptations for effectively learning from others, which include content-rich mechanisms that facilitate inferential reconstruction during cultural transmission as well as mechanisms that help learners select those members of their social world most likely to possess useful stuff to learn. Workshop members varied on how important or interesting these different elements were, but everyone seemed happy to get down to the business of sorting out when, where and how much. What we need is a large body of empirical work on specific cases.

4. Some of the important psychological mechanisms relevant for cultural evolution were not selected for improved socially learning, but for something else. Yet, they nevertheless influence the patterns and process of cultural evolution. The prevalence of bloodletting is a nice case example, as were some of the factors that influence Fijian food taboos.

5. Humans possess some improvisational intelligence, and this has important effects at least sometimes. However, there was disagreement about how important this intelligence is for explaining the world. There may also be differences on how important it is in human evolution.

6. The stability of cultural phenomena across generations is likely influenced by a rich multiplicity of factors, including teaching, learning biases, high fidelity copying, cognitive attractors, ecological constraints and socially-constructed learning environments (apprenticeships). It's simply now an empirical and theoretical question to understand when and where these are important and how they in-

fluence cultural evolution (micro-macro links). There was disagreement on the relative importance of teaching or pedagogy, but this can only be solved empirically on a case by case basis. No one thinks pedagogy is unimportant.

7. It's a worthy question to consider to what degree cultural evolution has driven genetic evolution. Some members, like me, would argue that cultural evolution emerged early in our lineage and has been a central driving force for over a million year. Others seem to favor the view that human evolution was driven by non-cultural factors, and cultural evolution arose later. Thus, while workshoppers had quite different priors on the relative importance, no one thought it was obviously stupid or a waste of scholarly effort.

8. Once terminological issues were clarified, most people agreed that it's a worthy line of investigation to consider how cultural group selection may have influenced cultural evolution. Of course, different workshoppers have quite different personal priors about how important this is likely to be, and this appeared to influence how they reacted to the empirical data. Nevertheless, no one was willing to publically defend the extreme and dismissive position taken by Steve Pinker in his EDGE essay.

On the flip side, I think there was some disagreement on the importance of spending a lot of time parsing terminology. It's my view that terms like "Darwinian" and "information" can be defined and deployed in a variety of ways, and we shouldn't care about the terms themselves. "Information", for example, seems to me to be used in the Dual Inheritance literature in a perfectly acceptable way (consistent with the use in engineering), just perhaps not the way some others have used the term (so what?). Arguably, "Darwinian" should be dropped entirely, since it can refer to populational processes in general or natural selection, or any kind of selection-like process. And, then, there's always the question of whether Darwinian refers to what Darwin actually thought vs. the neo-Darwinians. Rather than spending time arguing about whether culture is "Darwinian", I think time is better used going out and explaining some real world phenomena or building models that specify clear and identifiable processes.

THOUGHTS ON THE WORKSHOP (6)

By Olivier Morin

There seems to be a consensus (although I am sometimes not sure whether to include Blackmore in it) on several points:

— « Let's be Darwinian about Darwinism » (Dennett). Darwinian evolution should not be essentialized, there are vast grey areas between boundary cases and ideal Darwinian populations. In that spirit, we can make room for creative anticipation and transformation in our views of cultural evolution. Memes, like other scientific notions, can survive a radical re-thinking of the theory they first served to advertize: « I want to let the word "meme" go the way the words "atom" and "gene" went: de-Darwinize it! » (Dennett)

— Both transformation and selection drive cultural evolution. We dont know what their respective weights are (and the answer probably varies quite a lot from situation to situation) but there are reasons to think that these weights contribute to cultural evolution in an additive way: the less of one, the more of the other. Also, the proportion of transformation vs. replication determines what area we are in in Godfrey-Smith's space.

— We seem to agree upon the importance of distinguishing between recurrence and replication (Godfrey-Smith, Sperber). When a material is reliably transformed in a certain direction, it can reach an "attractor" state where any departures from its current state will cause it to be transformed back to the normal state (as Feher et al.'s experiment with finches illustrates). The kind of cultural transmission that follows can look a lot like replication (since similarity between models and descendants can be quite high), but the underlying mechanism is not replicative at all, since the parent's state does not cause its descendant's state. How much of the similarity we observe in cultural transmission is due to attraction, not replication, is not clear, but we should not underestimate the risks of wrongly positing replicative mechanisms when dealing with mere recurrence.

– One disagreement: Is evolution becoming more memetic, or less, or neither? Dennett argues that there has been a recent and ongoing de-Darwinization of culture. For Blackmore, on the other hand, « we are still in the primitive soup », cultural replication is in its infancy, and over time memes become better at replicating (as we see with cultural content that replicates not in our head, but in digital environments, with near-perfect fideliy — internet memes, "temes", etc.). Some others said they did not see a trend in either direction, without excluding the idea that cultures could move between different states in Godfrey-Smith's multi-dimensional space.

Cooperation and Cultural Group Selection

— Many versions of cultural group selection theory have been discussed, and there is no doubt that we agree on some of these versions. To recap, depending on the models, groups may be:

(1) Defined demographically, and competing by eliminating other groups (warfare, differential survival).

(2) Defined as political entities (i.e. sets of people cooperating together), and competing by attracting migrants from other groups. (For instance, if they could, most North Koreans would be fleeing to South Korea which has better institutions that prevent, for instance, mass starving. In this case, South Korea would not eliminate North Koreans as a population, but the Republic of South Korea as a polity could win over its Northern neighbour.)

(3) Defined by the use of certain institutions (which may be used independently by several independent political or demographic entities). Parliamentary regimes, for instance, have displaced other forms of government in many countries. Here, the competition is purely one of ideas: the set of all parliamentary regimes does not constitute a poitical unit (some parliamentary regimes are at war with other parliamentary regimes). Cultural forms simply compete against other cultural forms.

Against critics like Ruth Mace, who insist that the term "Group Selection" is not fit to cover such a wide range of definitions, Boyd insists that, from a modelling point of view, evolution in all these cases is fuelled by inter-group differences (with the caveat that "groups" mean quite different things in each of the three cases). There is broad agreement on the view that (at least the last two) mechanisms of "group selection" must play a crucial role in the history of cooperation for our species. There is indeed a disagreement on whether it is a good idea to lump together types 1, 2 and 3 of Cultural Group Selection, and to call them "Group Selection", given the wide range of diverse phenomena brought under this label, and the diversity of predictions yielded by each model.

— There is agreement on the importance of solving the problem of equilibrium selection. Reciprocity, punishment, etc. can (by virtue of the Folk Theorem) stabilize any interaction (cooperative, exploitative, etc.). We agree (I think) that cultural evolution is of great help in exploring the space of possible equilibria and select the best, as Henrich and others pointed out. One proposal that we did not discuss is the idea that equilibrium selection may be achieved through partner choice (with individuals

leaving the partners with whom they are stuck in a bad equilibrium) rather than through Cultural Group Selection. This proposal is quite popular in the literature on equilibrium selection. It is unclear, however, whether there would be any real difference between type-2 Cultural Group Selection and partner choice. (As noted, it is also controversial whether type-2 CGS constitutes group selection in a meaningful sense.)

— Group solidarity: (I don't know whether to count that one as a point of agreement or not. I was absent on Monday and the following is a reconstitution, mostly from breakfast conversations!) We take seriously the view that collecitve rituals can induce prosocial feelings, as stressed by Sterelny and Henrich, especially William McNeill's "dance and drill"; but the importance of the effects induced is still unclear. Also unclear is whether they can efficiently replace more banal cohesion-inducing forces, like coercion or material incentives, although Sterelny' discussion of Göbekli Tepe offered interesting arguments for this view. Sperber points out that we know very little about the actual workings of ritually-induced group cohesion, which must be treated as a kind of black box (as opposed to a truly naturalistic explanation).

Culture, adaptation and moves in Design Space

— There seems to be strong overall agreement on the adaptivity of culture once we reach the bottom line. As Boyd put it in his discussion of Blackmore, "Coevolution with a selfish entity does not in any way necessarily imply parasitism." (N.B. This is a good line to use in fights with one's significant other. I will make sure to re-use it ;o). The benefits of coevolution with culture, up till present times, haven't been much disputed, so I suppose we agree on them. There is less consensus regarding the present situation, and it seems obvious to all that some cultural innovations (at least) decrease human adaptation.

— The rise of complex and functional adaptation. There is complete agreement on the importance of cultural learning in creating complex norms or technologies, and no question that most such innovations would be beyond the reach of untaught evolved intuitions, a point rightly stressed by Henrich, Richerson, Boyd.

— Psychological biases (some of which are innate and widely shared in our species) influence the way we explore the Design Space of culture. Some forms are more likely to evolve due to what Richerson, Henrich & Boyd call "content biases". As Henrich points out, these biases may orient cultural evolution away from adaptive designs (Fiji food taboos). Sperber, Morin, Claidière argue that such biases may also have adaptive effects; obviously there is room for both claims to be true.

Some disagreement remains on related subsidiary issues:

(1) What is "improvisational intelligence" capable of? Both Dennett (in discussing "Bernstein's lament", the sadness of the composer who needs to manufacture a hit by painful trial and error) and RIcherson (along with Boyd and Henrich) strongly doubt that improvisational intelligence can by itself get us very far in Design Space. The truly important problems are simply too complex and too novel. I have argued that this may not always be true and I have presented several tentative examples (sound symbolism; pure coordination games; the design of writing systems) where individual intuitions seem to get it right without the contribution of a protracted and painful evolution by trial-and-error. What marks out these cases, I think, is that they are all cases where we play a game with other humans, not against nature. As Boyd et al. argue, it is extremely hard to anticipate what a good design for a canoe will be. Our intuitions are not completely helpless, they do restrict the Design Space in ways that are useful (few people try to build stone canoes)—but they only go so far. Yet, if the cultural inventions we talk about are meant to have their effects on other human minds, not on nature, our intuitions start to become much more helpful: it is much easier to ask one's intuitions what the ideal sound for the word "bad" would be, or how legible a given letter would turn out to be, etc. That is because we can instantly and effortlessly test these designs on our own brain.

(2) Are the respective contribution of cultural learning and individual cognition additive or interacting? Sperber, Morin (probably also Claidière) insist on the interplay between cultural transmission and evolved intuitions. They tend to take the view that cultural evolution will usually reinforce or magnify cognitive attraction. Henrich, Boyd and Richerson (it seems to me) often talk of cultural learning as an alternative to individual learning (rather than a complement). They talk of cultural influence (e.g. prestige or success biases) as a way of over-riding psychological biases. The respective strength of local cultural influences vs. general psychological constraints is difficult to measure if we do not know how both sets of factors come together (in an additive fashion or otherwise). Thus, the field is probably too young for there to be meaningful disagreements on this point.

THOUGHTS ON THE WORKSHOP (7)

By Peter Richerson

I have been thinking about the relationship between the kind of models that Pete, Joe and I have made and the ECM framework sketched by Dan and Nicolas in their paper. In particular, I have trying to understand how the mechanistic processes represented in our approach relate to homo-hetero distinction central to the ECM approach. In my remarks on Sue Blackmore's papers I sketched a simple model in which there were two variants with a selection stage and a transformation stage. So, first I'd like to convert this model to the ECM framework developed by Sperber and Claidière.

I do think that the disagreements among the various “schools” of cultural evolution represented at the meeting are relatively modest. I’ll first outline areas where I think disagreements are minimal and then raise some points where important issues may be outstanding.

Areas of broad agreement (taking it for granted that people will always disagree in detail)

1. Importance of cognitive processes. Dan S, Olivier, and Nico especially stress the myriad ways in which culture depends upon cognitive processes and in which cultural evolution is affected by such processes. They rightly stress that Rob, Joe and I have used simple models, as one necessarily must, to study the dynamics of cultural evolution. To my way of thinking, “theory” in fields like the evolutionary sciences consists of a toolkit of models, each itself fairly simple. We get at complex phenomena substantially by the piece-wise construction of families of models making different simplifying assumptions relevant to the specific scientific question at hand. Having models that represent cognitive processes more faithfully, perhaps by simplifying the population dynamic processes that Rob and I originally concentrated on, is work well worth doing. Joe has branched out in that direction. Rob and his students have a Bayesian learning model unifies the individual and social learning inference process. Many studies of social learning in humans and animals provide a lot of data on the one-generation-to-the next time scale that one might use to test such models. For example, I think that Olivier’s flop problem is quite real and is worthy of formalizing.

2. The concept of attraction. Until my reading for this meeting I did not appreciate how broad a set of phenomena attraction represents. I did not appreciate that local environmental contingencies could act as attractors. For Rob and I, and I think Joe, such contingencies are critical for assembling intricate cultural adaptations fairly quickly and in turn key for understanding how a costly cognitive apparatus for managing culture might have evolved in the first place. We may still have some different guesses about the importance of different classes of attractors but this is chiefly an empirical matter which is likely to be settled in due course.

3. The diversity of cases. Peter GS and Dan D make a very good case that some examples of both organic and culture are more paradigmatically "Darwinian" than others. (I use the scare quotes to mark that the paradigm is not Darwin's own formulations but the mid 20thCentury Neo-Darwinian Synthesis with its hard gene based notions of transmission, the rigid proximate-ultimate distinction, and other things foreign to Darwin.) I have no problem with this idea. Rob's and my old models contained parameters that we imagined were under selection that measured the strength of faithful transmission versus the strength of, essentially, attractors. We paid disproportionate attention to the case where the transmission effects are fairly strong relative to attractors on the per generation time scale because the evidence suggested to us that humans are unusual in this regard.

4. I don't have a big problem with the concept of memes so long as the meme-gene analogy is not excessively rigid. Susan assures is that Rob's, Joe's and my old fears in this regard are unfounded.

Perhaps larger issues still outstanding

1. I continue to be impressed with the ongoing cognitive and comparative work on social learning. Kim's "evolved apprentice" seems to be exactly what humans are. The work of people like Susan Carey, Elizabeth Spelke, Paul Harris, Mike Tomasello, Karen Wynn, Paul Bloom, and their students and colleagues have constructed a reasonably detailed picture of how development in infancy brings on line the cognitive machinery that makes humans a much more imitative species than any other studied to date. In typical animal social learning, the rate and fidelity of transmission is relatively low and the role of attractors is necessarily much stronger. Comparative studies of the aptitude of chimpanzees versus children for social learning highlight the unique importance of imitation and perhaps pedagogy in humans. For many culturally transmitted traits—syntax, word meanings, artifact construction skills, social norms and institutions—fidelity of transmission itself is very high. Biases and guided variation do modify and sort among variants acquired by social learners and that is very important too, especially to the extent that it sensitive to environmental reinforcement. On Rob's, Joe's and my account, such transmission fidelity coupled to even weakish environmental reinforcement is necessary to understand the evolution of complex cultural adaptations and the spectacular cultural adaptive radiations of humans. From an ecological point of view, humans are the analog of thousands of biological species each carrying a mixture of arbitrary historical contingent differences and often exquisite adaptations to local circumstances. Why we came to have the capacity to do this and how it works are the biggest questions for evolutionary study of humans. Culture has worked so spectacularly for

us you might think that any number of other lineages would have stumbled across high fidelity culture over the last 600 million years since complex animals have existed. I'm not sure everyone agrees with this agenda.

2. I have qualms about the concept of universal Darwinism. Culture/memes are a lot like genes in some respects and not like them in others. I see a shallow analogy where others seem to see a fundamental law-like similarity. Analogies are really useful things. They allow you to borrow concepts and models from other fields and save yourself having to invent too much new stuff. But the disanalogies are important too. At some point you've squeezed all the useful work out of the analogy and pursuing it any further will lead you down false paths. As Rob mentioned in the meeting, the population genetics style of modeling is a disciplined but flexible framework for milking the meme-gene analogy dry and then setting off on new paths dictated by the phenomenology of the problem you are interested in. For example, Donald Campbell, and Gerald Edelman suggested that the population genetic approach should be applicable to cognitive development. Olivier's word models and Dan S' model of attraction might be tackled with such a framework. Shrink the notional time step from a generation to a day in the life of a child learning new words from others that it interacts with. Children exercise some repertoire, hopefully a small repertoire, of attractors to acquire words others, who may or may not engage in a little pedagogy. They actively use some words and build a larger passive vocabulary, day by day. Thinking shallow analogies honors the diversity of processes that we call organic and cultural evolution and cognition.

3. Multi-level selection and multi-modal selection. Mark Pagel and Susan give primacy to selection on genes and selection on memes respectively. Rob convinced me early on in our partnership that sweeping generalizations about what "selection" favors are dangerous. Natural selection on genes admits to a number of modes. Selection on simple linear fitness gradients does one thing. Allow for rough fitness topographies and selection produces different results (as Peter G-S notes). Throw in density and frequency dependent selection. Add antagonistic selection in males or females, or at different ages. Mate choice and artificial selection introduce agent based rather than natural selection, demigod designers if you want. With cultural evolution agent based social selection runs wild. The institutions of a society set up systems of rewards and punishments that heavily impact genetic fitness. Polly Wiessner once argued to me that in the Ju/'hoansi (!Kung San) that there is little punishment even for murder. But a multiply homicidal male is felt to be dangerous to the camp. If his relatives are willing to take the responsibility, and they normally are, he is semi-exiled to a distant camp where he lives with his relatives. At best, he can attract only an old woman as a wife. His genetic fitness is zeroed out as much as if he were executed.

Similarly, it is easy to imagine selection acting at multiple levels of organization, especially on cultural variation. On the one hand, horizontal transmission of cultural variants sets up the possibility of selection for selfish memes. At the same time horizontal transmission is adapted to spread individual or

group functional innovations to a large population rapidly. Genetically selected attractors tend to act like an immune system, favoring genetically advantageous cultural variants, as Mark P argues.

Institutions tend to be systems characteristic of a whole society or substantial parts of it. Such variation is an easy target for group selection. Some evolutionists want the selective design process to be a relatively orderly maximization-of-fitness process. True, absent that, the idea that selection produces adaptations is problematical. But I think that many evolutionary biologists consider the various forms of selection to result in a rather anarchic process. Cultural evolution, being a recently evolved system, is rather less orderly than the genetic system. Anarchy is as much the rule as the exception. In the Holocene, cultural group selection seems to have ever larger and better integrated social systems.

Against this large scale trend, every lineage of large scale societies seems to be a series of hot house flowers that bloom and bust. Despite invitations in Susan's and Mark's material and Rob's and my presentation on cultural group selection we did not have deep conversation at Santa Fe. I sense that we would be pretty various on the issues of modes and levels of "selection."

THOUGHTS ON THE WORKSHOP (8)

By Dan Sperber

Let me fist express my heartfelt gratitude to Dan for this great initiative, and to Louis Godbout and to the SFI for making it possible. It has been a wonderful workshop of serious, demanding, insightful, informal, friendly discussion of a kind and quality rarely experienced.

I would like to thank all the participants for their contributions, which have all been inspiring, and also for their willingness to entertain and help develop the idea of cultural attraction that has been a long time in the making, but that is still very much in a work-in-progress stage. I am particularly grateful to Pete and Rob and Joe, who have done so much more for so long to develop our understanding of cultural evolution and our capacity to model it. Given the sheer qualitative and quantitative importance of their work in the area and their unique level of expertise, they could have been, if not dismissive, at least much more severe in their reaction to the work of the attraction gang. Instead of which, they have been attentive, constructive and really very helpful. This is what I had hoped for, but was not sure of getting. Rob's own post-meeting written comments shows how, given their expertise, they are in a position to improve on our suggestions and to make the very idea of attraction a better articulated one, especially but not uniquely on the modelling side, and I hope they will.

While much of the workshop turned on the differences and convergences between the two main approaches represented, I want to say also that I got a lot from Dan, Kim, Peter, and Sue, both from their own ideas –I would have liked to have had time to discuss Kim's and Peter's ideas in particular in much greater detail — and from their contributions to the framing of the exchanges. I am sorry, in the same vein that Mark could not make it. It would have been great to have had his viewpoint. Hopefully a next time.

Several of us have already listed points of agreement and differences, and I have nothing much to add to these summaries (that's an unfair benefit of being the last one). I agree with Dan – and I guess everybody else that a follow up, possibly enlarged, would be a good idea. At this stage, I will just share two or three of the several ideas that emerged from our discussions and that I would like to articulate and discuss in greater depth.

Adaptive and non-adaptive aspects of culture, and how to explain them. Cultures are full of adaptations or at least adaptive traits — and this is crucial to explaining the overall success (so far) of humans – but they are also of items that prima-facie are not or hardly adaptive.

Regarding cultural adaptive traits (with examples such as knots and canoes and guns and computers, but also efficient institutions) one issue is how much do they owe their existence and evolution to people's understanding of their usefulness and effective efforts at improving them, as opposed to selective forces without foresight? This is, I take it we agree, a more or less rather than an either-or question, with different answers for different items.

Take cases where people's understanding, foresight, and inventiveness are quite determinant. Does this mean that, in those cases at least, we have to change model and move to rational choice theory or something of the sort? None of us believe this. For me, in particular, not if we take into account attraction. Mental processes that can be assessed as 'rational' are just psychological factors of attraction among other. If presented with two variants of a tool, one of which is more efficient than the other, the second variant may be, because if greater efficiency is recognizable with human cognitive capacities, an attractor just as is the English pronunciation of "data" relative to the Latin pronunciation in an English-speaking population. Take a tool that is efficient but for which there is a close-by possible but not yet actual even more efficient variant. Cognitive capacities can be a factor that make this not yet present improved version attractive with a big homo-attraction potential that would kick in one the variant has been instantiated. And so on. The point is that adding attraction to the cultural evolution story allows to integrate evolved mechanisms that tend to produce rational choices, not as an alternative kind of explanation, but as a factor of attraction among many.

Take now cases of cultural traits that look non-adaptive, as there are many in religion, in the arts, and so on. One strong tendency in traditional social sciences and in evolutionary approaches alike is to try and show that they are adaptive (or 'functional') after all. With some genuine successes, actually. But, I would argue, with much overshooting, and a general tendency to look for the benefits and underestimate the costs. In general anyhow, the adaptationist explanations of religion, music, and so on, however good they may be, have little or nothing to say about the details of religious beliefs and rituals, the evolution of musical traditions, and so on, which are, of course — and for excellent scholarly reasons — of prime interest to social scientists, cultural anthropologists, historians, and so on. Factors such as 'prestige bias' may explain why such traits persist but, by themselves make no predictions about their specific contents. To the extent that these traits are not arbitrary – and I believe they rarely are – a variety of factors of attraction, some pretty general, other historically contingent, should at least greatly help do the job. Note that aesthetic judgment (and other psychological dispositions that are not or at least not clearly a matter of rationality) as psychological factors of attraction are not, from an explanatory point of view, that different from efficiency judgments.

Opacity: This relates to the previous point. At one point during a break, Rob and I agreed that we agreed that a most crucial feature of cultural evolution was the opacity to the people of much of the contents that they acquire and propagate. Indeed. This comes up everywhere, in technology, in religion, in norms and institution, and so on. People only partly understand what they are doing or thinking and why they are doing or thinking it. This opacity – which is a matter of degree of course – is what makes social transmission so important. It plays, I believe, a crucial role in the acceptability of cultural traits: it is, in important ways easier to trust what you don't fully understand and hence cannot properly evaluate on its own merits. The work of Gergely and Csibra on the role of opacity in natural pedagogy and on the role of natural pedagogy in cultural transmission comes also to mind.

Memes and memetics: There have been great insights in Dawkins' whole idea of memes even if it failed to spawn a successful scientific program. The idea that a trait that causes its own propagation will be successful is compelling and forces one to rethink many generally accepted idea on cultural traits. A trait may help cause its own propagation because it favors the reproductive success of its carriers, but this is in no way a necessary condition. Or it may compromise its propagation by killing off its carriers, but this is rare. Many if not most cultural traits – with some blatant exceptions –, in their relatively short historical lives – again, there are blatant exception, especially on the side of technology — do not have effects on the fitness of their carriers strong enough to either benefit or harm themselves greatly. Note that the rightfully challenging 'meme's eye view' is not compromised if a meme is redefined as an attractor (or replaced by the notion of an attractor). Attractors or meme qua attractors don't even need to benefit themselves, let alone their carriers, they need to be in a position to benefit from the pool of cultural traits in which they occur. In fact, I would argue – but not here – the challenge then becomes even more interesting.

THOUGHTS ON THE WORKSHOP (9)

By Kim Sterelny

1. Attraction and selection. What is the relationship between the SMC attractor based framework, and the BRH approach, which prima facie finds a more obvious role for selection and adaptation, especially cumulative adaptation? While the SMC group do not deny the existence of examples of adaptive complexity, their discussion and models are not organized around these kinds of cases. In the discussion, I tried out one way of seeing how the two approaches might fit together, borrowing from the saner versions of evo-devo: see the mechanisms identified in the SMC approach as constraints and biases in the supply of variation to selection-like processes. Except in extreme cases where those constraints reduce the supply of variation to a trickle (as in C's zebra finch case, we the experimental set up exposes the hatchling to a single model) that can still leave plenty of work for selection-like forces. I take it the SMC suggestion was to see the attractor-approach as more general than selectionist approaches, which for them come out as a special form of attraction. That seems to me to obscure an important distinction between the supply of variation and its fate. But maybe M's suggestion does not fit that evo-devo picture, seeing he focuses on downstream constraints that affect success rate (on whether a cultural variant is a flop). But what is the difference between that and selection (selection often depends on factors internal to a lineage, as in the relationship between sexual selection and sensory biases).

2. How "blind" is cultural transmission/cultural learning? Or better: how important are relatively blind processes, where agents adopt new cultural variants (a) without actually understanding the effects of these variants on their lives, and why they have those effects; or (b) without even taking themselves to understand those effects (perhaps by some form of unconscious imitation; perhaps because the agent acquiring the CV acts on some form of father-knows-best principle. I take it that the selectionist formal models abstract away from this issue: a selectionist model of bow-improvement or kayak making will apply just fine, even if each incremental improvement in technology is guided by causal understanding of why the old design worked, and why the new design will work just a little better. But it seems to me that when the transmission and improvement process is relatively blind, the cultural se-

lectionist models are explanatory in a way that they are not, if the incremental changes are made and adopted because the individual agents understand what is gong on (obviously, though, blindness is a matter of degree).

Perhaps one of the differences between the SMC group and the BRH group is that the SCM group seem to emphasize somewhat more the role of the agent's own resources to the transmission process.

3. Imitation. A closely related issue is the role of imitation and its nature (what are its cognitive requirements; again, what is the role of causal understanding in the transmission process). I have been skeptical about the role of imitation in the cultural transmission of (at least) technique, artisanship, (perhaps until relatively recent composite technologies). Rob Boyd thinks I am flat wrong, massively overselling the role of "field testing" technique as an agent acquires a capacity (so he thinks imitation is both way important, and not especially demanding of causal understanding: see 2 above). I remain utterly unpersuaded. Dan Sperber helpfully pointed out (and I carry on about this in The Evolved Apprentice, too) that in the transmission of complex skills (i) demonstration is very important, and (ii) demonstration is not the input to imitation. Moreover, (iii) demonstration depends pretty clearly on some reasonable causal understanding of the capacity in question, since it often involves a meta-commentary as well as stylized/exaggerated action patterns.

4. Tokens, types, memes. It seems pretty clear that when we focus on tokens of cultural variants: (i) typically, any given token in a transmission network will have multiple parents; (ii) these often play unequal roles in the transmission process; (iii) jointly, (i) and (ii) above will make it theoretically difficult and/or computationally intractable to specify the fitness of cultural tokens. Prima facie, if we cannot identify parent offspring relations, and hence cannot define either fitness or hereditary, we cannot apply the Darwinian framework.

Can we think of selection as acting on types? After all: it is often true that cultural variants are increasing (of decreasing) in relative frequency within a culture, and that that change is plausibly caused by the effects of cultural variants on those that adopt them. One bow design (perhaps using fletched arrows) is displacing another in a community, at a time, because it is fitter. Up to a point, I think this is OK, but talk of fitness here is only serving as a form of score-keeping; it is just a measure of the size of the ensemble of tokens; fitness will not explain anything about the size of that ensemble — the explanation will come from a comparison of the mechanics of the two bows; the ecological and economic upshot of those differences, and the various social and psychological mechanisms through which those bow-design-differences became salient to individual. Cultural fitness becomes a purely statistical notion. Admittedly, there is a line of thought that says that about biological fitness too (but not one I accept).