

Partial connectivity increases cultural accumulation within groups

Maxime Derex^{a,1} and Robert Boyd^b

^aInstitute of Human Origins, Arizona State University, Tempe, AZ 85287; and ^bSchool of Human Evolution and Social Change, Arizona State University, Tempe, AZ 85287

Edited by Adam T. Powell, Max Planck Institute for the Science of Human History, Jena, Germany, and accepted by the Editorial Board February 2, 2016 (received for review September 21, 2015)

Complex technologies used in most human societies are beyond the inventive capacities of individuals. Instead, they result from a cumulative process in which innovations are gradually added to existing cultural traits across many generations. Recent work suggests that a population's ability to develop complex technologies is positively affected by its size and connectedness. Here, we present a simple computer-based experiment that compares the accumulation of innovations by fully and partially connected groups of the same size in a complex fitness landscape. We find that the propensity to learn from successful individuals drastically reduces cultural diversity within fully connected groups. In comparison, partially connected groups produce more diverse solutions, and this diversity allows them to develop complex solutions that are never produced in fully connected groups. These results suggest that explanations of ancestral patterns of cultural complexity may need to consider levels of population fragmentation and interaction patterns between partially isolated groups.

cultural evolution | innovation | population size | social network | technological trajectory

People everywhere rely on technology for their survival (1). In even the simplest foraging societies, essential tools are beyond the inventive capacities of individuals; they result from a cumulative process by which innovations are gradually added to existing cultural traits across many generations (2). Recent work suggests that a population's ability to develop complex technologies is positively affected by its size and connectedness (3–5). Large interaction networks allow individuals to learn from many others, and theory predicts that increased opportunities to learn socially reduces the rate of cultural loss and increases the rate at which people improve existing cultural traits. This prediction is supported by evidence from both field and laboratory studies (5–10). However, some authors have been reluctant to embrace this idea, pointing out discrepancies between measures of population size and observed cultural complexity (11–14). It seems likely that factors other than social network size affect the evolution of cultural complexity, and these factors may obscure the effect of the population size under some conditions.

One possibility is that social network structure matters as well. Economists have recently investigated how levels of connectedness in networks affect a group's ability to solve problems of varying complexity (15, 16). In agreement with previous work in cultural evolution, they found that high levels of connectivity help groups to solve simple tasks. However, they also suggest that well-connected networks perform poorly at solving complex tasks (15, 16). The complexity of a problem can be represented by mapping all possible solutions onto a measure of performance, which results in what has been called a fitness landscape (17). Simple problems are associated with smooth landscapes, each with a unique optimum that will be reached by local exploration from any point in the landscape. By contrast, complex problems are associated with rugged landscapes with multiple peaks of different heights. In rugged fitness landscapes, a propensity to learn from successful models can cause the entire population to converge rapidly on a suboptimal

peak (18). Thus, a well-connected group will quickly reach a local optimum that may not be the highest one.

In the context of cultural evolution, this work suggests that well-connected populations might not exhibit the most complex cultural repertoires. Technologies typically arise from a cumulative process that operates through incremental changes within path-dependent technological trajectories and by combining traits that have evolved along different trajectories (19–21). Screws, for example, are relatively simple artifacts that can be improved incrementally by using better materials, modifying head shape, or using a different kind of thread. However, one may also combine screws with unrelated cultural traits to produce radically new technologies. The bench vise, for instance, results from the combination of a screw with a lever, and the wheel barrow combines a lever with a wheel (22). More recently, combinations of levers, pulleys, cranks, ropes, and toothed gears resulted in the production of early machines that were used for milling grains, irrigation, construction, and time-keeping (23). The production of such innovations strongly depends on cultural diversity because more cultural traits provide more combinatorial opportunities. Thus, well-connected populations may be less likely to produce complex technologies because the ability to learn from the most successful models can reduce fitness landscape exploration and cultural diversity.

To address this question, we used an experimental task in which individuals had to discover successive innovations to produce a virtual remedy and stop the spread of a virus. To make the process of cultural accumulation realistic, we specified that innovations were contingent upon earlier discoveries and resulted from incremental

Significance

The remarkable ecological success of the human species has been attributed to our capacity to overcome environmental challenges through the development of complex technologies. Complex technologies are typically beyond the inventive capacities of individuals and result from a population process by which innovations are gradually added to existing cultural traits across many generations. Recent work suggests that a population's ability to develop technologies is positively affected by its size and connectedness. Here, we present an experiment demonstrating that partially connected groups produce more diverse and complex cultural traits than fully connected groups. This result suggests that changes in patterns of interaction between human groups may have created propitious conditions for the emergence of complex cultural repertoires in our evolutionary past.

Author contributions: M.D. designed research; M.D. performed research; M.D. and R.B. analyzed data; and M.D. and R.B. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission. A.T.P. is a guest editor invited by the Editorial Board.

¹To whom correspondence should be addressed. Email: maxime.derex@gmail.com.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1518798113/-DCSupplemental.

improvement or recombination of different traits (19, 24). Players were provided with six active ingredients that could be associated in groups of three to create a remedy. Whereas all 56 possible triads provided players with a measure of remedy efficiency, two triads provided players with a new active ingredient (A_1 and B_1 ; Fig. 1). Each of these new ingredients allowed players to create new, more rewarding triads simulating incremental improvements. The opportunity for players to discover two different innovations from the same initial set of ingredients simulated cultural divergence. By allowing players to produce specific triads, early innovation events opened alternative path-dependent trajectories: Triads using A_1 allowed players to discover the new ingredient A_2 , whereas B_1 -based triads allowed players to discover B_2 . The same principle applied to the discovery of third-level ingredients, after which players had to combine innovations from both A and B path-dependent trajectories to progress further in the landscape. For example, the discovery of one of the two fourth-level innovations required triads involving both A_3 and B_3 .

Participants were placed within groups of six. Each participant had 72 trials to refine his/her own remedy and maximize his/her cumulative score. We compared two different treatments. In the fully connected condition, players were provided with information about the remedies produced by the five other group members after each trial. In the partially connected condition, players were

members of a partially connected network of three subgroups of two players. Players could observe the solutions of the other player within their subgroup, and, occasionally, different groups were connected by the movement of individuals between groups.

We predicted that individuals in fully connected groups would tend to converge on the same solution, making them unable to generate new traits by combining solutions from different path-dependent trajectories. We expected individuals in partially connected groups to be more likely to progress along different path-dependent trajectories, and thus exhibit more diverse cultural repertoires, including innovations resulting from the combination of traits arising from different evolutionary pathways.

Results

We found that fully connected groups were able to discover A_3 (50% of groups) and B_3 (33.3%), but none of the fully connected groups were able to discover both A_3 and B_3 . In comparison, 91.7% of partially isolated groups discovered A_3 , 66.7% discovered B_3 , and 58.3% discovered both (Fig. 2). Additionally, every partially connected group that reached both third-level innovations combined them together and produced some fourth-level innovations. In three of the five partially connected groups that failed to reach fourth-level innovations, the three subgroups happened to produce innovations along the same pathway, thereby

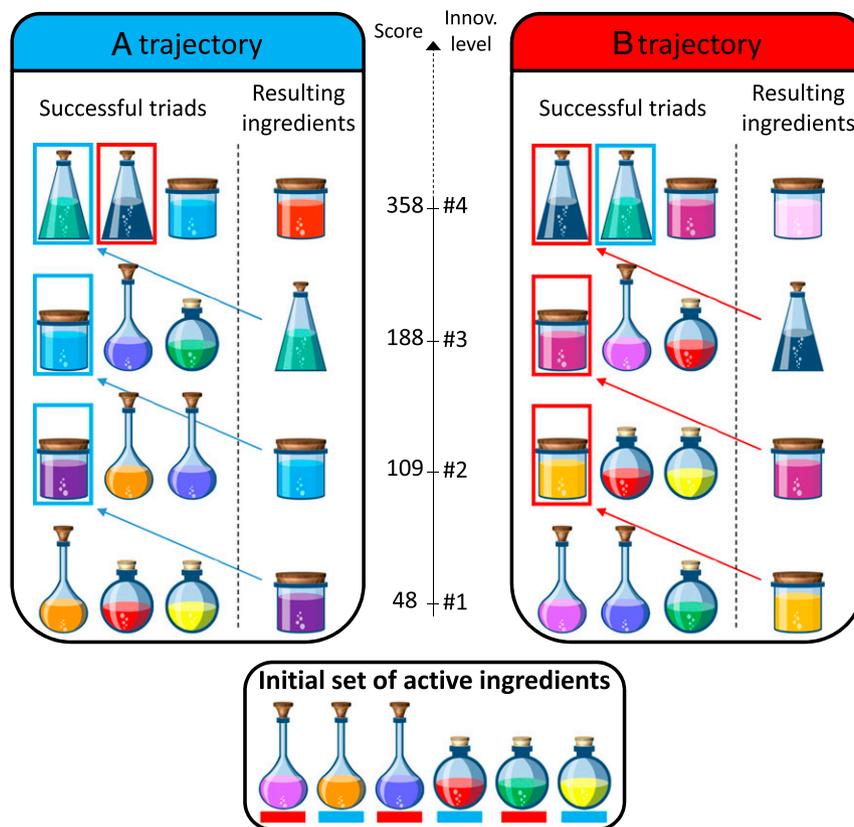


Fig. 1. Path-dependent trajectories. The six initial ingredients were randomly divided into two types unknown to players. These ingredients could be associated using a three-slot apparatus. Within each type, ingredients were randomly assigned one of three possible values. Ingredients positively interacted with ingredients from the same type and negatively with ingredients from the other type, which created an initial fitness landscape with two possible optima with the same value. The discovery of one of these optima provided players with a new ingredient (A_1 or B_1). A_1 and B_1 were assigned the same value (superior to those values associated with the initial ingredients) and positively or negatively interacted with the same number of randomly chosen ingredients (the interaction pattern of A_1 was the reverse of the interaction pattern of B_1). As a result, A_1 and B_1 allowed players to produce specific and more rewarding triads (A_1 -based and B_1 -based triads). Scores associated with superior-level innovations were attributed according to the same principle, and the most rewarding A_1 -based and B_1 -based triads provided players with a new ingredient (A_{i+1} and B_{i+1} , respectively). A_3 and B_3 , however, were assigned specific properties and strongly and positively interacted with each other. As a result, the most rewarding A_3 -based triad (that provided players with A_4) required B_3 , and, reciprocally, the most rewarding B_3 -based triad (that provided players with B_4) required A_3 .

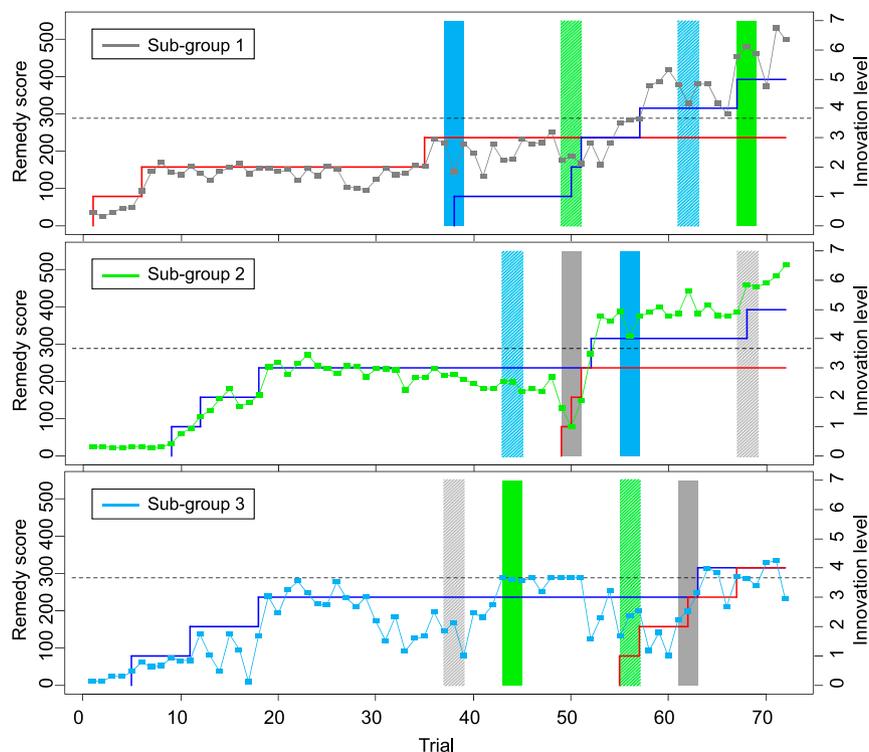


Fig. 4. Example of partially connected group scores and associated progress along trajectories. Subgroup 1 initially progressed along the B trajectory (red line), whereas subgroups 2 and 3 progressed along the A trajectory (dark blue line). Subgroup 1 benefited from the visit of a subgroup 3 member (vertical blue bar) and acquired A_1 , before reaching A_3 after visiting subgroup 2 (vertical green hatched bar). Subgroup 2 reached B_3 following the same visit, whereas the first contact with subgroup 3 (vertical blue hatched bar) did not result in any improvements for either subgroup because they previously produced the same innovations. Subgroups 1 and 2 independently reached A_4 following these first contacts. Subgroup 3 eventually got out of the local optimum after second contacts with subgroups 1 and 2. Note that after contacts, subgroups were still progressing along different trajectories as subgroups 1 and 2 exploited A_5 and subgroup 3 exploited B_4 . Vertical full and hatched bars illustrate incoming and outgoing visits, respectively, by a single individual. Bar color illustrates subgroups with whom contacts took place. The horizontal black dotted line illustrates the best score that can be reached when exploiting a single evolutionary trajectory.

hard to disentangle these effects properly. However, assuming that long-distance flows of raw materials and personal ornaments can provide information about degree of population structuring, it should be possible to investigate the extent to which changes in the pattern of interaction among members of different populations affected cultural complexity in our evolutionary past.

The effect of connectedness may explain discrepancies between inferred population size and observed cultural complexity. For example, Powell et al. (5) used molecular data to estimate when different regions of the world reached the same population density as Europe at the start of the Upper Paleolithic. They report that the crossing of the density threshold coincided with the appearance of markers of modern behavior in some regions, but not others. Our results suggest that population size alone is not sufficient to predict cultural complexity. We predict that populations with more markers of group identity should exhibit more complex cultural traits than populations of the same size that have fewer markers of group identity.

Our results seem to be at odds with the results of a previous experimental study of problem solving, which found that more connected groups were better able to solve a task associated with a rugged landscape than less connected ones (30). However, this study was based on a fitness landscape that was very different from the fitness landscape assumed in our experiment. In the task used by Mason and Watts (30), participants searched for the most rewarding position on a 2D map, and every landscape position was accessible from every other position in the landscape. Even if groups got trapped on a suboptimal peak, single nonlocal moves could reveal more rewarding solutions. Although this kind of

landscape is likely to apply to some optimization problems, it does not capture the cumulative nature of the technological process. Historians of technology and economists have described how innovations are contingent upon earlier discoveries (19, 24), which means that innovations create new, more rewarding and previously unreachable positions in the landscape. Our experiment captures this dynamic because as soon as groups start along a path-dependent trajectory, nonlocal moves become unprofitable.

Our experiment illustrates how networks of partially isolated groups connected by occasional migration events can outperform fully connected networks of the same size. It is, however, worth noting that other network structures may affect exploration and cultural diversity in a similar way. Theoretical work by economists, for example, typically compares networks composed of permanent links with different levels of clustering rather than occasionally connected groups. These studies indicate that more clustering leads to more thorough exploration of the design space, which suggests that high levels of isolation might not be necessary to allow connected groups to evolve on different pathways. In theory, fragmentation should only play a role when the rate at which innovations spread to other groups is lower than the rate at which groups produce innovations. Weakly connected groups, for example, may not progress along different trajectories if the rate of evolution is low because innovations will spread among groups before alternative solutions are produced. We saw this case in our experiment. In the partially connected treatment, two groups had subgroups that initially progressed along different trajectories but then converged to the same path-dependent trajectory after contact

events revealed alternative, more rewarding solutions to less efficient subgroups. As a result, these whole groups converged to a single local optimum and failed to reach both third-level innovations.

Cultural homogenization may also depend on other factors, such as the structure of cultural learning. For instance, individuals are much more likely to learn skills from kin and members of their communities than from others. Contemporary data from Fijian villages show that being from the same village doubles an individual's likelihood of being selected as a model (31). Such preferences constrain information flow and, according to recent experimental results, may facilitate the emergence of different cultural norms (32). More generally, conformism (the tendency to acquire the most common behavior exhibited in a group) can also help maintain significant differences among groups despite factors such as migration and intermarriage (33). In addition, large and geographically widespread interconnected populations may experience and take advantage of diverse local resources, which could make them more likely to explore different solutions. Heterogeneity in stone-flaking systems observed in southern Africa during Marine Isotope stage 5 (130–80 ka), for example, is thought to have resulted from a low rate of information transfer between groups whose technological systems were strongly locally adapted (27).

Further research will be needed to clarify the extent to which patterns of between-group interactions or population structure may have affected cultural complexity in our evolutionary past. In particular, theoretical work could investigate the effect of variation in the properties of the fitness landscape; population structure; and structure of cultural learning, such as payoff-biased learning, conformism, and accessibility of cultural models. Recent research indicates that contemporary hunter-gatherer societies display a unique social structure involving extensive interactions between people living in different residential groups (34, 35). This population structure may constrain information dissemination and promote exploration of the design space. However, high interaction rates probably prevent bands within an ethnolinguistic group from progressing along radically different technological trajectories. Instead, the effect we report is more likely to be a result of contacts between different ethnolinguistic groups. Archeologists are just starting to infer ancestral population structure from the archeological record (27, 28, 36). Our results suggest that it is worth pursuing this effort. Taking into account levels of population fragmentation and between-group contact may shed new light on ancestral patterns of cultural accumulation.

Methods

Participants. A total of 144 University of Montpellier students (72 women and 72 men) were randomly selected from a database managed by the Laboratory of Experimental Economics of Montpellier (LEEM) and recruited by email from various universities in Montpellier, France. Informed consent was obtained from all subjects before starting the experiment (ethical approval was given by the Arizona State University Institutional Review Board, code: STUDY00002815). The subjects ranged in age from 18 to 44 y (mean of 24 y, SD of 4.32 y). Participants received €5 for participating and an additional amount ranging from €5 to €30 depending on their own performance.

Procedure. The experiment took place in a computer room at the LEEM at the University of Montpellier. For each session, a maximum of 18 participants (exclusively male or female) were recruited and randomly assigned to one condition of the experiment. Participants sat at physically separated and networked computers, and were randomly assigned to a group. Players did not know who belonged to their group and were instructed that communication and note taking were not allowed. Before starting the experiment, participants were requested to enter their age and sex, and could read instructions on their screens. At the end of the game, each subject received a reward according to his/her performance (€15 on average).

Game Principle. The participants played a computer game (programmed in Object Pascal with Delphi 6) in which they were asked to develop a remedy to fight a virtual virus. Players were initially provided with six basic active

ingredients that could be used without any limit and could be associated in groups of three to create a remedy. All triads were allowed, including those trials involving the repeated use of the same ingredient. The order of the ingredients had no effect on the result, so that 56 unique triads could be produced from the six initial ingredients. Whereas all triads provided players with a score (as a measure of remedy efficiency), two of them allowed players to benefit from new active ingredients. New ingredients arose when players produced a triad that belonged to a list of predetermined successful triads. When discovered, new ingredients could, in turn, be associated with other ingredients. Triads using new ingredients allowed players to produce more rewarding triads and created opportunities to find other new ingredients. A total of 16 new ingredients could be produced. Players were given 25 s to generate a triad and were asked to maximize their cumulative score across 72 trials. No information associated with triads (score or resulting ingredient) was displayed before the end of 25 s.

Fitness Landscape. The fitness landscape associated with our task was designed to allow players to progress along two symmetrical path-dependent trajectories. To do so, we randomly divided the initial set of ingredients into two types unknown to players. Within each of the types, ingredients were randomly assigned one of three possible values (6, 8, or 10). Triad scores based on the initial set of ingredients were calculated as follows:

$$\text{Score} = \frac{(1 + 0.5\alpha) \cdot (S_1 + S_2 + S_3)}{\beta}, \quad [1]$$

with α taking the value 0, 1, or 2 depending on whether triads involved one, two, or three different ingredients and β taking the value 1 or 2 depending on whether triads involved ingredients from one or two types. S_1 , S_2 , and S_3 are the scores of ingredients 1, 2, and 3, respectively. As a result, two different triads based on initial ingredients provided players with the highest payoff (Fig. 1). The discovery of these triads provided players with a new ingredient that allowed them to produce new and more rewarding triads (Fig. S1).

New ingredients were given a score that was equal to the score of the best initial ingredients (10) plus $5 \times i$, with i equal to the innovation rank ($A_1 = 15$, $A_2 = 20$, etc.). Then, ingredients were randomly divided into two types, so that ingredients previously positively interacting with each other did not necessarily do so when they were associated with a new ingredient. Triad scores were then calculated according to Eq. 1, except that 90% of the score of the best possible triad from the lower level was added to the result. For example, B_2 -based triad scores were calculated according to Eq. 1, and 90% of the best B_1 -based triad was then added to the score. This bonus ensured that scores of lower level innovation-based triads were only slightly overlapping with triads involving higher innovations (Fig. S1). Each time a player generated the most rewarding triad, given his/her set of ingredients, he/she was provided with a new ingredient.

To simulate innovations that result from the combinations of traits arising from different evolutionary trajectories, we assigned two pairs of ingredients specific interacting properties (A_3/B_3 and A_6/B_6) that led to higher payoffs. The score of triads containing one of these pairs was calculated according to Eq. 1, and 150% of the score of the best possible triad from the lower level was added to the result. As a result, fourth- and seventh-level innovations could only be produced by combining A_3 with B_3 and A_6 with B_6 , respectively. Again, this bonus ensured that this form of innovation (combining traits coming from different pathways) provided players with scores that did not overlap with triads involving single trajectory-based triads (Fig. S1).

Treatments. We compared two different treatments that involved same-sized groups of players but provided them with different social learning opportunities. In the fully connected condition, players were part of a group of six and were provided with information from their other group members after each trial (the same five sources of social information). In the partially connected condition, players were part of a network of three subgroups of two players, and were subsequently connected by the movement of individuals between groups. During the first 36 trials, players could only observe the other member of their own subgroup. At the end of trial 36, one player was randomly removed from his/her own subgroup and joined another subgroup for three trials, after which he/she returned to his/her initial subgroup for three trials. Then, another contact event took place. Each individual was moved from his/her initial subgroup to another subgroup once. In total, each subgroup experienced four contact events, two incoming and two outgoing, with two different subgroups (Fig. 4). All treatments involved 72 participants in single-sex populations (12 replicates each).

Social Information. After each trial, players were provided with information about triads that they produced and triads produced by other individuals with whom they were connected [ingredients that had been used, triad score, and resulting new ingredient (if any)]. This information was displayed for 5 s, after which players were provided with an opportunity to produce a new triad. Although the social information panel was removed from the screen, players could obtain access to this information through the use of a record panel. The record panel provided players with their last triad score, the ingredients involved, and a record of their new ingredients (if any). Players could click onto new ingredients to get a reminder about how to produce them. By clicking onto an anonymized name (e.g., “player 3”) and associated last triad score, players could switch between their own record and the record of players they were connected with. The record panel for other players displayed the player’s last triad score, ingredients involved, and a record of their new ingredients (if any). Players could learn how to produce these ingredients by clicking onto them.

Tutorial and Pregame Information. Before starting, the players had to complete a tutorial during which basic actions, such as dragging and dropping ingredients, had to be completed. The tutorial also guided players’ actions

until they could access (nonrelevant) social information to make sure that all players mastered the game interface before starting the experiment. Players were informed that the ultimate aim of the game was to maximize their cumulative score and that new ingredients were generally more efficient than initial ones. Players were also informed that their monetary reward depended on their cumulative score. The fitness function that determined the value of a triad was unknown to players.

ACKNOWLEDGMENTS. We thank all the people who helped test the game, Dimitri Dubois for his help during experimental sessions, and the Laboratory of Experimental Economics of Montpellier (LEEM, University of Montpellier) for hosting the experiment. We also thank Kim Hill, Curtis Marean, Sarah Mathew, Charles Perreault, Joan Silk, Deborah Strumsky, Jose Lobo, and people from the Evolutionary Foundations of Human Uniqueness project for helpful discussions. Last, we thank three anonymous reviewers for their constructive comments that helped to improve this paper. This research was made possible through the support of Grant 48952 from the John Templeton Foundation to the Institute of Human Origins at Arizona State University. The opinions expressed in this publication are those of the authors and do not necessarily reflect the views of the John Templeton Foundation.

- Richerson PJ, Boyd R (2005) *Not by Genes Alone* (Univ of Chicago Press, Chicago).
- Boyd R, Richerson PJ, Henrich J (2011) The cultural niche: Why social learning is essential for human adaptation. *Proc Natl Acad Sci USA* 108(Suppl 2):10918–10925.
- Shennan S (2001) Demography and cultural innovation: A model and its implications for the emergence of modern human culture. *Cambridge Archaeological Journal* 11(1):5–16.
- Henrich J (2004) Demography and cultural evolution: How adaptive cultural processes can produce maladaptive losses—the Tasmanian case. *Am Antiq* 69(2):197–214.
- Powell A, Shennan S, Thomas MG (2009) Late Pleistocene demography and the appearance of modern human behavior. *Science* 324(5932):1298–1301.
- Kline MA, Boyd R (2010) Population size predicts technological complexity in Oceania. *Proc Biol Sci* 277(1693):2559–2564.
- Derech M, Beugin M-P, Godelle B, Raymond M (2013) Experimental evidence for the influence of group size on cultural complexity. *Nature* 503(7476):389–391.
- Derech M, Boyd R (2015) The foundations of the human cultural niche. *Nat Commun* 6: 8398.
- Muthukrishna M, Shulman BW, Vasilescu V, Henrich J (2013) Sociality influences cultural complexity. *Proc Biol Sci* 281(1774):20132511.
- Kempe M, Mesoudi A (2014) An experimental demonstration of the effect of group size on cultural accumulation. *Evol Hum Behav* 35(4):285–290.
- Collard M, Buchanan B, O’Brien M (2013) Population size as an explanation for patterns in the Paleolithic archaeological record: More caution is needed. *Curr Anthropol* 54(S8):S388–S396.
- Buchanan B, O’Brien M, Collard M (2015) Drivers of technological richness in pre-historic Texas: An archaeological test of the population size and environmental risk hypotheses. *Archaeol Anthropol Sci*, 10.1007/s12520-015-0245-4.
- Collard M, Buchanan B, O’Brien MJ, Scholnick J (2013) Risk, mobility or population size? Drivers of technological richness among contact-period western North American hunter-gatherers. *Philos Trans R Soc Lond B Biol Sci* 368(1630):20120412.
- Read D (2012) *Population Size Does Not Predict Artifact Complexity: Analysis of Data from Tasmania, Arctic Hunter-Gatherers, and Oceania Fishing Groups* (UCLA Center for Human Complex Systems, Los Angeles).
- Fang C, Lee J, Schilling MA (2009) Balancing exploration and exploitation through structural design: The isolation of subgroups and organizational learning. *Organization Science* 21(3):625–642.
- Lazer D, Friedman A (2007) The network structure of exploration and exploitation. *Adm Sci Q* 52(4):667–694.
- Wright S (1932) The roles of mutation, inbreeding, crossbreeding and selection in evolution. *Proceedings of the Sixth International Congress of Genetics* 1:356–366.
- Mason WA, Jones A, Goldstone RL (2008) Propagation of innovations in networked groups. *J Exp Psychol Gen* 137(3):422–433.
- Dosi G (1982) Technological paradigms and technological trajectories. *Res Policy* 11(3):147–162.
- Fleming L (2001) Recombinant uncertainty in technological search. *Manage Sci* 47(1): 117–132.
- Youn H, Strumsky D, Bettencourt LMA, Lobo J (2015) Invention as a combinatorial process: Evidence from US patents. *J R Soc Interface* 12(106):20150272.
- Wagner A, Rosen W (2014) Spaces of the possible: Universal Darwinism and the wall between technological and biological innovation. *J R Soc Interface* 11(97):20131190.
- Arthur WB (2009) *The Nature of Technology: What It Is and How It Evolves* (Simon and Schuster, New York).
- Basalla G (1988) *The Evolution of Technology* (Cambridge Univ Press, Cambridge, UK).
- David P (2007) Path dependence: A foundational concept for historical social science. *Cliometrica* 1(2):91–114.
- Martin R, Sunley P (2010) The place of path dependence in an evolutionary perspective on the economic landscape. *Handbook of Evolutionary Economic Geography*, eds Boschma R, Martin R (Edward Elgar, Cheltenham, UK), pp 62–92.
- Mackay A, Stewart BA, Chase BM (2014) Coalescence and fragmentation in the late Pleistocene archaeology of southernmost Africa. *J Hum Evol* 72:26–51.
- Vanhaeren M, d’Errico F (2006) Aurignacian ethno-linguistic geography of Europe revealed by personal ornaments. *J Archaeol Sci* 33(8):1105–1128.
- Bar-Yosef O (2002) The Upper Paleolithic revolution. *Annu Rev Anthropol* 31: 363–393.
- Mason W, Watts DJ (2012) Collaborative learning in networks. *Proc Natl Acad Sci USA* 109(3):764–769.
- Henrich J, Broesch J (2011) On the nature of cultural transmission networks: Evidence from Fijian villages for adaptive learning biases. *Philos Trans R Soc Lond B Biol Sci* 366(1567):1139–1148.
- Centola D, Baronchelli A (2015) The spontaneous emergence of conventions: An experimental study of cultural evolution. *Proc Natl Acad Sci USA* 112(7):1989–1994.
- Henrich J, Boyd R (1998) The evolution of conformist transmission and the emergence of between-group differences. *Evol Hum Behav* 19(4):215–241.
- Hill KR, et al. (2011) Co-residence patterns in hunter-gatherer societies show unique human social structure. *Science* 331(6022):1286–1289.
- Hill KR, Wood BM, Baggio J, Hurtado AM, Boyd RT (2014) Hunter-gatherer inter-band interaction rates: Implications for cumulative culture. *PLoS One* 9(7):e102806.
- Kovacevic M, Shennan S, Vanhaeren M, d’Errico F, Thomas M (2015) Simulating geographical variation in material culture: Were early modern humans in Europe ethnically structured? *Learning Strategies and Cultural Evolution During the Palaeolithic, Replacement of Neanderthals by Modern Humans Series*, eds Mesoudi A, Aoki K (Springer, Tokyo), pp 103–120.