

Know your network: people infer cultural drift from network structure, and expect collaborating with more distant experts to improve innovation, but collaborating with network-neighbors to improve memory

Emory Richardson (emory.richardson@yale.edu),
Daniela Miro-Rivera (daniela.miro-rivera@yale.edu), & Frank Keil (frank.keil@yale.edu)
Yale Department of Psychology, 2 Hillhouse Ave. New Haven, CT 06511 USA

Abstract

We suggest that some of the mechanisms underlying network effects on cultural evolution are intuitively accessible to laypeople, and may be part of the suite of social learning strategies underlying the human capacity for cumulative culture. Interest in the psychological mechanisms underlying this capacity typically focuses on learners' ability to identify reliable sources and capacity for high-fidelity imitation. Yet, at the population level, research suggests that network structures themselves may influence cumulative learning by changing individuals' explore-exploit patterns. In our experiments, adults infer that more proximal or distal clusters in a fragmented network will have more similar or dissimilar technological "styles", and prefer to seek advice from more distant experts when asked to innovate, but more proximate experts when asked to remember. Commonsense intuitions about how social networks shape our access to information and diversity-fidelity tradeoffs for memory and innovation may make us more effective social learners.

Keywords: cumulative culture; innovation; intuitive theories; networks

Introduction

Isaac Newton is sometimes praised for recognizing the cumulative nature of discovery and innovation in pointing out that only "standing on the shoulders of giants" had enabled him to see further (or more deeply) than others. However, models of cumulative cultural evolution suggest that even the emphasis on "giants" may still be misplaced; innovation appears to be driven not by rare individual geniuses, but by a community of knowledge which accumulates and recombines diverse elements over time, to produce incremental advances that may then be further recombined to produce more incremental advances in the future (Mesoudi & Thornton, 2018; Derex & Mesoudi, 2020). Contra Newton, the remarkable pace of technological development in human societies may depend less on the stature of its giants and more on commonsense intuitions about which kinds of ideas can be fruitfully recombined, and which otherwise-average people are likely to have access to those kinds of ideas.

Of course, this is not to deny that expert judgment is typically more valuable than non-expert judgment. Indeed, innovation (as opposed to rediscovery) may require an increasingly high degree of expertise as technological knowledge accumulates, potentially spanning multiple fields (Hardwig, 1991). However, while expert knowledge may improve a learners' judgment about where innovation is

possible, likely, or needed, innovation critically depends on the discovery of something previously unknown rather than familiarity with existing knowledge. And expertise does not necessarily produce innovation.

Rather than explaining innovation as the spark of individual genius, evidence from studies of cultural evolution suggests that diversity is key: repeatedly recombining diverse elements to produce incremental advances can lead to innovative breakthroughs without relying on genius *ex machina*. Importantly, the diversity of the elements available for a population to recombine can be increased or decreased by manipulating the size and the structure of a social network. Though a community can simply be too small to maintain a broad knowledge base (Kempe & Mesoudi, 2014; Derex et al., 2013), subdividing sufficiently large social networks into smaller clusters can increase diversity by changing explore-exploit decisions (Derex & Boyd, 2016). Though conformist tendencies may still influence individual agents' exploration patterns *within* clusters, *between*-cluster influences are reduced, allowing clusters to drift apart. Restoring the lines of communication between clusters then allows them to combine what they've learned. In "rugged" fitness landscapes, which contain multiple good-but-not-optimal solutions, fragmenting networks can thus increase learning by encouraging individual learners to explore more diverse options (Mason, Jones, & Goldstone, 2008). Simulation studies and in-lab experiments suggest that these manipulations can dramatically increase the speed of cultural accumulation. Moreover, in at least some domains technological improvements can accumulate over time even if individual agents have no understanding of the causal mechanisms underlying the technology they're developing (Derex, Bonnefon, Boyd, & Mesoudi, 2019). Thus, innovation may not only have no need for individual geniuses, but could simply be a slow but inevitable result of social learning in networks, with individual understanding improving as a largely separate process.

However, evidence of sophisticated and early-emerging intuitions about "who knows what" in our social networks (for review, see Harris, Koenig, Corriveau, & Jaswal, 2018) suggests an alternative: even if we lack any extraordinary capacity for individual genius, our capacity for cumulative culture could be accelerated by using social learning strategies that combine knowledge of "who knows what" with (A) a tendency to seek out more or less diverse perspectives as appropriate, and (B) intuitions about where

to find such perspectives. For instance, a recent study of “hot streaks” in scientific and artistic careers found that while individuals focus on a narrow topic during the hot streak itself, streaks are typically preceded by a period of exploration of diverse styles and topics (Liu, et al., 2021). Similarly, publications in the sciences produced by authors who had not previously collaborated with each other are more “disruptive” and have greater multi-disciplinary impact than publications by frequent collaborators; moreover, these effects are stronger for first-time collaborators who are more distant from each other in their networks (Zeng, et al., 2021). We suggest that these are not simply abstruse academic descriptions of the idiosyncratic behavior of geniuses; rather, they reflect commonsense strategies for learning from your social network.

For individuals to benefit from exploring diverse disciplines or seeking out diverse collaborators, some understanding of both (A) how different domains of knowledge cluster together and (B) how knowledge spreads through our social networks would seem to be necessary. For instance, a chemist who wants to solve a problem in chemistry will clearly benefit less from exploring a diverse range of painting techniques than a diverse range of topics more closely related to chemistry, but staying within one’s own cloistered subfield may cut the flow of fresh perspectives to a trickle (Aral & Alstynne, 2011). Even 5 year olds recognize that knowledge clusters into domains of expertise, and seek out sources with domain-relevant expertise over sources whose expertise is not domain-relevant (Lutz & Keil, 2002). Here, we ask whether information seeking may similarly be guided by commonsense intuitions about how knowledge spreads through social networks and the contingent benefits of seeking out diverse sources of knowledge, and develop a procedure adaptable to developmental work, with a goal of conducting similar studies with children.

Experiment 1a

Participants in Experiments 1a-b were introduced to an avatar who lived on an island (Fig. 1) with two “societies” living in distinct social networks, depicted as silhouettes with black lines connecting the people who “talk with each other the most often”, and separated by a mountain range in the middle of the island. Participants were told that the people on each side of the island were expert boat-makers, who learned boat-making skills from their parents and other adults on their side of the island. One of the boat-makers (“Max”) needed to ask another expert for help. Participants were told that they would be shown different experts Max could ask, and after hearing his question, they would decide which expert would give more helpful advice to Max. Participants were also asked to rate how similar they believed each boat-making expert’s “style” would be to Max’s style.

In Exp. 1a, we asked whether participants would infer that experts more closely connected to Max in the network would have more similar boat building styles than experts further away, and whether participants would infer that

while proximal experts would better help Max remember specific techniques, distant experts would better help Max innovate new techniques.

Participants. We recruited 80 participants from MTurk for Experiment 1; an additional 23 participants were screened out prior to participating for twice failing to answer three basic comprehension questions about the instructions. Participants were assigned to one of two conditions: 38 participated in the Memory condition, and 42 in the Innovation condition.

Procedure. The experts were presented in pairs (Fig 1); in three trials, the physical distance between the experts was equal (as measured by gridlines on the map which participants were told indicated physical distances on the map), but one expert was in an entirely different network (e.g., the experts in the Southwest and Southeast villages in Fig 1 are physically equidistant from Max, but while there are six nodes separating Max from the Southwest expert, there is no network connection between Max and the Southeast expert). In three more trials, the pairings were crossed, so that one expert was either more proximal than the other in only network distance (degrees of separation), only physical distance, or in both. In the Innovation condition, participants were told that no one on either side of the island had ever been able to build a boat that could travel more than 100 miles, and that Max needed help figuring out how to build this new kind of boat; in the Memory condition, participants were told that Max needed help remembering how to tie a specific kind of knot that he needed for his boat. Advice ratings and Similarity ratings were given on a 10-point scale. However, for Similarity, participants inferred how similar each expert’s “style” was to Max; for Advice, participants rated which of the two experts would be more helpful to Max. Finally, because the order of the Advice and Similarity ratings was counterbalanced, participants who completed the Similarity questions first did not yet know what question Max would be asking the expert. As no order effects were found in any experiment, counterbalancing will not be discussed further.

Results. We first asked whether participants inferred that experts closer to Max in the network would have more similar building styles. Network distance was undefined for the “Global” experts (from the other side of the island), but each Global expert was paired with a “Local” expert (from Max’s side of the island) an equal physical distance from Max; and, Local experts differed amongst themselves in their degrees of separation from Max. Thus, we created a numeric dummy variable matching each Local expert with the physically distant Global expert (NW vs. NE, MW vs. ME, SW vs. SE). A linear mixed effects model with random intercepts by participant, and LocalVsGlobal and SourceMatch as fixed effects (Fig. 2) suggested that participants expected the building style of Local experts to share a high degree of similarity with Max’s ($\beta_{Int} = 9.99$, $SE = 0.34$, $p < .001$), but significantly less similarity with Global experts ($\beta_{Global} = -6.11$, $SE = 0.42$, $p < .001$). Moreover, participants expected Local experts further from Max to have less similar building styles than those closer to




Figure 1. Social networks in Exps. 1a-b and Exp. 2. Participants rate *how similar* each source’s building style will be to Max’s, and *which* of two sources will be *more helpful* (Exps. 1a-b) or *how helpful* each source will be (Exp 2) for memory and for innovation. In Exp. 1a, contrasting NWvsNE, MWvsME, and SWvsSE controls for physical distance; NWvsSW, NWvsSE, SWvsNE contrast network & physical distance. In Exp. 1b, contrasts examine cultural drift RedGreen, RedPurple, GreenPurple, PinkNavy, SkyOrange. Exp. 2 uses naturally occurring social structures (cities & schools).

him ($\beta_{SourceMatch} = -1.04, SE=0.14, p < .001$), but with an interaction suggesting that this effect was weaker for Global experts ($\beta_{Inxn} = 0.74, SE = 0.20, p < .001$), who varied amongst themselves only in physical distance, but (unlike Local experts) did not vary in network distance (which was undefined for all Global experts insofar as the two networks were not connected).

We next asked if participants also believe that the experts’ optimal distance from Max depends on whether he’s asking for help remembering or innovating. We analyzed participant ratings using a linear mixed effects model with random intercepts for each participant and QuestionType as a fixed effect; in order to compare Advice ratings to chance for each condition, we deleted the intercept and centered Advice ratings on the midpoint of the scale (5), such that a higher rating indicates more help from the more distant expert. We found that participants in the Innovation condition expected more distant experts to be significantly more helpful than more proximal experts ($\beta_{Innovate} = 2.33, SE = 0.35, p < .001$), while participants in the Memory condition expected proximal experts to be significantly more helpful than distant experts ($\beta_{Memory} = -3.77, SE = 0.37, p < .001$).

Notably, physical distance appeared to have no impact on participants’ helpfulness judgments, even in the three trials in which the relevant expert (out-of-network for Innovation, in-network for Memory) was physically further away. However, inferences about which expert is likely to be more helpful may differ from inferences about which expert is worth seeking help from; indeed, both children and adults expect others to weigh value of information against the physical cost of acquiring it (Aboody, Zhou, & Jara-Ettinger, 2021; Baker et al., 2017). Future work could examine how learners weigh the costs and benefits of learning from their social networks as compared to learning from their physical environments. For instance, it may be easier to access the friend-of-a-friend-of-an-expert than the expert themselves; yet, the extent to which third- or fourth-

hand information retains (or gains) value for the learner may depend on how much learners expect it to be distorted (or refined) in transmission.

Finally, we asked whether the inferred similarity between Max’s building style and the experts’ was related to the experts’ inferred helpfulness. Because participants had rated the similarity of each expert’s building style to Max’s, but rated each expert’s helpfulness relative to another expert, we computed a relative similarity variable for each of the expert pairs. However, while participants believed that experts with more similar boat-building styles to Max’s would better help him remember a specific knot for his boat ($\beta_{SimDiff} = -0.15, SE = 0.05, p < .001$), participants did not infer that experts with less similar styles would better help Max invent a new kind of boat ($\beta_{SimDiff} = 0.06, SE = 0.05, p = .20$).

Experiment 1b

In Exp. 1b, we link the two networks, allowing us to extend Exp. 1a in two ways. First, connecting the two networks provides a more sensitive measure by allowing us to compute the absolute and relative distances between sources, instead of only contrasting “Local” and “Global” sources as in Exp. 1a. We predicted that (A) people would infer that sources more distant in the network would have building styles less similar to Max’s, and that (B) as the *relative* distance between two sources increased, so would the degree to which participants would see the more distant source as more helpful to Max in innovating (Exp. 1b does not include a Memory condition simply because our primary goal was to examine the effect of network distance on innovation). Second, Exp. 1b contrasts the “linking” sources themselves, allowing us to ask whether participants infer a “cultural drift” effect in addition to a network distance effect: if so, one more degree of network distance should make a greater difference if it crosses a network boundary than if not, because sources are less influenced by networks to which they are less closely connected.

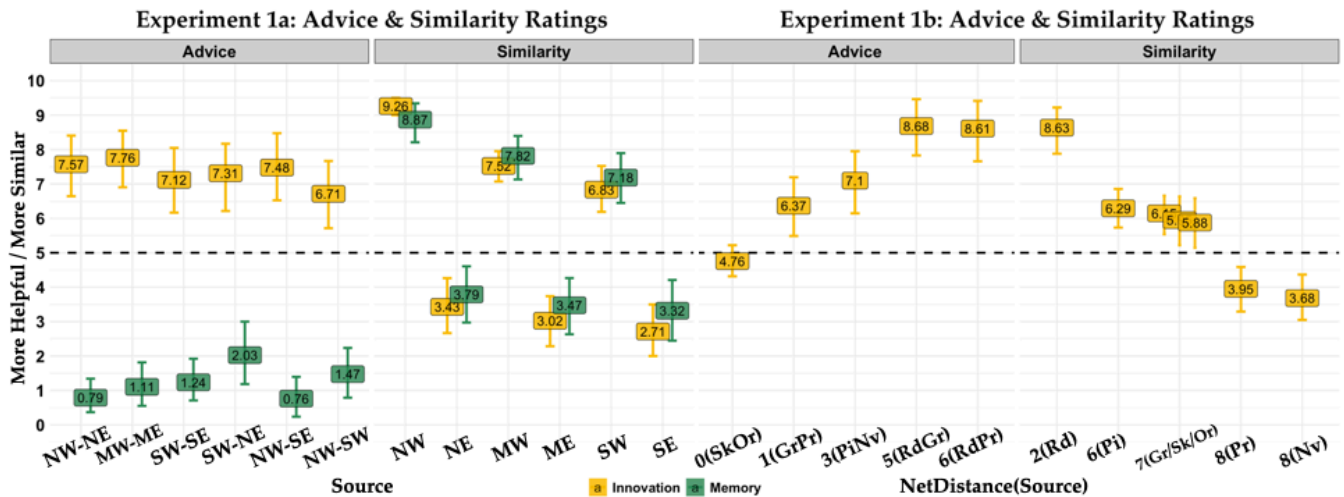


Figure 2. Means and 95% CIs for Exps. 1a-b. No Memory condition was run in Exp. 1b. **Advice:** which source (distal vs. proximal) would be *more* helpful. **Similarity:** how similar would given source be to Max. In Exp. 1b, sources are arranged by network distance (Similarity: *absolute* degrees of separation from Max. Advice: *difference* in degrees of separation from Max).

Participants. We recruited 41 participants from MTurk for Experiment 1b; an additional 2 participants were screened out prior to participating for twice failing to answer three basic comprehension questions about the instructions.

Procedure. The procedure in Experiment 1b was the same as Experiment 1a except for changes in the network structure and the position of the expert sources in the network itself. As in Experiment 1a, Advice ratings and Similarity ratings used 10-point scales, and the order of the blocks was counterbalanced.

Results. We first asked whether participants inferred that experts closer to Max in the network would have more similar building styles. Because the networks on each side of the island were linked (unlike in Experiment 1a), we computed network distance directly as the degree of separation between Max and each expert. A linear mixed effects model with random intercepts by participant, network distance as a fixed effect, and similarity ratings centered on the midpoint of the scale (5) suggested that while participants expected the most proximal experts' building styles to be similar to Max's ($\beta_{(Int)} = 3.63, SE = 0.34, p < .001$), they also expected less similarity in building styles from experts increasingly distant from Max in the network ($\beta_{NetDist6} = -2.34, SE = 0.39, p < .001$; $\beta_{NetDist7} = -2.64, SE = 0.32, p < .001$; $\beta_{NetDist8} = -4.68, SE = 0.39, p < .001$; $\beta_{NetDist9} = -4.95, SE = 0.39, p < .001$). Moreover, comparing experts on either side of the network "bridge" linking the two sides of the island suggests that participants may infer a cultural drift effect: while similarity ratings did not differ for the experts highlighted (in Pink and Green or for the experts highlighted in Purple and Navy ($\beta_{PinkGreen} = 0.15, SE = 0.24, p = .54$; $\beta_{PurpleNavy} = 0.27, SE = 0.24, p = .26$), participants did infer significantly greater similarity to Max's building style from the expert highlighted in Green than the one in Purple ($\beta_{GreenPurple} =$

$2.20, SE = 0.24, p < .001$). In other words, while one additional degree of separation from Max made no difference to participants if the experts were on the same side of the "bridge" between the two networks, crossing the bridge did appear to make a difference.

We next asked if participants also believe that more distant experts will be more helpful to Max for innovation questions. We analyzed participant ratings using a linear mixed effects model with random intercepts for each participant and the difference in network distance between the two experts as a fixed effect; in order to contrast Advice ratings with chance, we centered Advice ratings on the midpoint of the scale (5). When the two experts were an equal distance from Max, participants did not expect one to be more helpful than the other ($\beta_{(Int)} = -0.24, SE = 0.42, p = .56$); however, participants increasingly expected the more distant expert to better help Max innovate as the relative distance between the experts increased ($\beta_{Net1} = 1.61, SE = 0.45, p < .001$; $\beta_{Net3} = 2.34, SE = 0.45, p < .001$; $\beta_{Net5} = 3.93, SE = 0.45, p < .001$; $\beta_{Net6} = 3.85, SE = 0.45, p < .001$).

Finally, we asked whether the inferred similarity between Max's building style and the experts' was related to the experts' inferred helpfulness. Because participants rated the similarity of each expert's building style to Max's on an absolute scale, but had rated which of two experts would be more helpful given the kind of help Max needed on a relative scale, we computed a relative similarity difference for each of the expert pairs to compare with participants' relative helpfulness ratings. Participants believed that experts with more dissimilar boat-building styles to Max's would better help Max invent a new kind of boat ($\beta_{SimDiff} = 0.38, SE = .06, p < .001$).

Experiment 2

In Experiments 1a-b, the network structures were drawn out for participants explicitly. However, while there is

evidence that even young children represent their social networks (e.g., their school) in considerable detail and with very high accuracy (Gest, Farmer, Cairns, & Xie, 2003; Capella, Neal & Sahu, 2012), it is unlikely that these representations are as explicit as we make them by drawing out the networks on physical maps. Still, naturally occurring social structures may serve as reasonable proxies for network distance by suggesting how many mutual connections two individuals are likely to share. For instance, even without knowing whether any two agents know each other directly, it's reasonable to assume that the probability that two agents know each other is higher if they are, e.g., randomly selected from the same school than from the same city, or from the same city than from the same country. Indeed, recent work suggests that children infer that schoolmates are more likely to know the rules of their school than close friends from other schools, but friends are more likely to share personal secrets — even if the friends are from a different country (Lieberman, Gerdin, Kinzler, & Shaw, 2020).

Thus, in Experiment 2, we manipulated network distance implicitly; participants were told about a state engineering tournament in which teams first competed within-class, then within-school, within-city, and finally between-cities. If participants infer that teams are influenced most by those they communicate with most often, this implicit manipulation of network structure may lead them to make similar inferences about seeking assistance from in-network or out-of-network sources for memory and innovation as in Experiments 1a-1b. Moreover, this manipulation may be more feasible for use with children in later work.

Participants. We planned to recruit 60 participants from MTurk for Experiment 1b; however, due to a coding error, 1 participant who should have been screened out for twice failing to answer three basic comprehension questions about the instructions was nevertheless able to complete the study. They are excluded from the analyses below; however, including them changes nothing in the results. Participants were assigned to one of two conditions: Memory (n=29) or Innovation (n=30).

Procedure. Participants were introduced to a protagonist “Max”, whose class was participating in a drone-building tournament in a “State Science Club”, along with other science classes from their own city and one other city. The structure of the tournament was described as follows: teams consisted of pairs of students from the same science class who competed in weekly contests; during the first month, teams competed each week against the other teams from their class; during the second month, teams competed each week against all of the other teams from their class and other classes in their school; during the third month, contests were between all of the teams from all of the schools in the same city. Finally, during the last month, all of the teams from both cities were to compete in a final contest. However, during the final month each team could choose one other team to work with in a group, from any of the teams in either city. In the Memory condition, participants were told that during the first week, Max’s

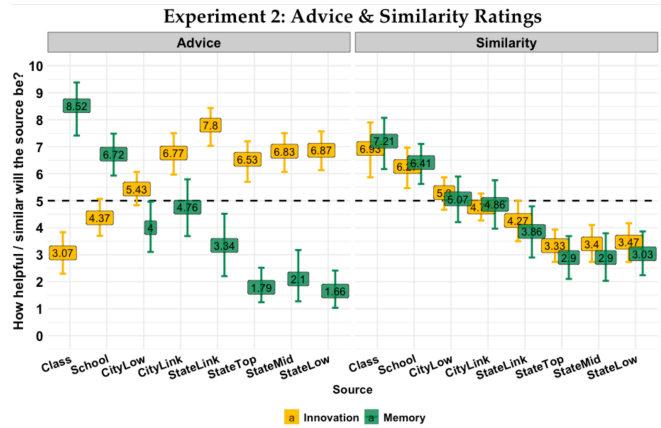


Figure 3. Means and 95% CIs for Experiment 2.

drone had used a special kind of propeller to do a special trick; before the final contest, Max and his partner realized that the trick they had done would let them win the contest — but, they couldn’t remember which propeller was needed to do the trick. In the Innovation condition, Max and his partner realized that to win the contest, they would need to build a new and creative kind of drone, unlike what any of the teams had ever built before. Finally, participants were told that while “most of the students in Westlake don’t know any of the students from Eastview, because they only see each other once a year at the state finals...there are two students from Westlake Science Club who each have a close friend in the Eastview Science Club, so they talk about Science Club all the time”; thus, comparably to Experiment 1b, Experiment 2 includes two “bridges” directly linking the two networks.

As in Experiments 1a-1b, participants were shown a set of expert sources (here, each source was a team, selected for their implicit network distance from Max’s team), and asked to rate how similar each team’s building style would be to Max’s, and how helpful each team would be in winning the contest. As in Experiments 1a-b, the order of the Advice and Similarity measures was counterbalanced; however, unlike in Experiments 1a-b, participants in Experiment 2 rated how helpful each team would be on an absolute 0-10 scale (0 meaning “not helpful at all” and 10 meaning “very helpful”) instead of a relative 0-10 scale (where 0 and 10 meant that the near expert or far expert was “definitely” more helpful, respectively). Moreover, it should be noted that while in Experiments 1a-b participants might assume that Max had never spoken to the more distant sources (though strictly speaking, the network visualizations were only said to show who spoke “most often”), in Experiment 2, the teams were explicitly said to have all competed against one another at some point during the tournament. Thus, while in Experiments 1a-1b, the influence of more distant sources in the network was implicitly suggested to be only indirect, in Experiment 2 it was explicitly direct (though less in degree: more distant teams competed against each other fewer times, and presumably shared fewer mutual connections because of pre-existing social structures intrinsic to class-school-city hierarchies).

Results. Analogously to Experiments 1a-1b, we first asked whether participants inferred that more “in-network” teams would have more similar building styles to Max’s team. Because the network distance was implicit, we analyzed similarity ratings in two ways: by dividing teams into “Local” (from Max’s city) and “Global” sources (from the other city), and by creating a dummy variable ordering the teams by implicit distance from Max’s team (Class < School < CityLow < CityLink < StateLink = StateTop = StateMid = StateLow). A linear mixed effects model with random intercepts by participant and LocalVsGlobal as a fixed effect suggested that participants expected Local teams’ building styles to be more similar to Max’s than Global teams’ ($\beta_{LvGGlobal} = -2.45, SE = 0.18, p < .001$). Similarly, a linear mixed effects model with random intercepts by participant, implicit network distance as a fixed effect, and similarity ratings centered on the midpoint of the scale (5) suggested that while participants expected the most proximal experts’ building styles to be similar to Max’s ($\beta_{(Int)} = 2.07, SE = 0.29, p < .001$), they also expected less similarity in building styles from experts increasingly distant from Max in the network ($\beta_{NetDist2} = -0.73, SE = 0.34, p < .03; \beta_{NetDist3} = -1.88, SE = 0.34, p < .001; \beta_{NetDist4} = -2.25, SE = 0.34, p < .001; \beta_{NetDist5} = -3.00, SE = 0.34, p < .001; \beta_{NetDist6} = -3.89, SE = 0.28, p < .001$).

Next, we compared the helpfulness ~ network distance relationship across conditions. As with similarity ratings, we analyzed advice-helpfulness ratings both by categorizing teams as either Local and Global and by using a dummy variable to arrange the teams in order of their implicit distance from Max’s team; both models included random intercepts for each participant and QuestionType as a second fixed effect. In the Innovation condition, participants’ beliefs about how helpful the Local teams would be did not differ from the midpoint of the scale, ($\beta_{(Int)} = -0.09, SE = 0.29, p = .75$), but they believed Global teams would be significantly more helpful than Local teams ($\beta_{LvGLocal} = 2.10, SE = 0.30, p < .001$). Participants in the Memory condition were significantly more likely than participants in the Innovation condition to believe that Local teams would be helpful ($\beta_{Memory} = 1.09, SE = 0.41, p < .009$), but believed that Global teams would be significantly less helpful for Memory questions than Local teams ($\beta_{QType*LvG} = -5.88, SE = 0.43, p < .001$). Results were more granular when network distance was treated as a continuous variable: while participants in the Innovation condition expected the team in Max’s own class to be unhelpful overall ($\beta_{(Int)Innovate} = -1.93, SE = 0.42, p < .001$), they expected teams to be more helpful with each additional degree of distance from Max’s own class ($\beta_{NetDist2} = 1.30, SE = 0.52, p = .012; \beta_{NetDist3} = 2.34, SE = 0.52, p < .001; \beta_{NetDist4} = 3.70, SE = 0.52, p < .001; \beta_{NetDist5} = 4.73, SE = 0.52, p < .001; \beta_{NetDist6} = 3.68, SE = 0.42, p < .001$). Conversely, participants in the Memory condition expected the team in Max’s own class to be significantly more helpful than participants in the Innovation condition ($\beta_{Memory} = 5.45, SE = 0.60, p < .001$), and they expected teams to be less helpful with each additional degree of distance from Max’s own class

($\beta_{NetDist2*Mem} = -3.09, SE = 0.74, p < .001; \beta_{NetDist3*Mem} = -6.88, SE = 0.74, p < .001; \beta_{NetDist4*Mem} = -7.46, SE = 0.74, p < .001; \beta_{NetDist5*Mem} = -9.91, SE = 0.74, p < .001; \beta_{NetDist6*Mem} = -10.34, SE = 0.60, p < .001$).

Finally, we examined the advice ~ similarity relationship across conditions. A linear mixed effects model with random intercepts by participant, with Similarity*Condition as fixed effects suggests that participants in the Innovation condition believed that the less similar a team’s building style was to Max’s, the more helpful they would be in inventing an entirely new kind of drone ($\beta_{Similar} = -0.50, SE = 0.07, p < .001$). Conversely, participants in the Memory condition were less sure than participants in the Innovation condition that teams with the most dissimilar building styles would be helpful to Max ($\beta_{Memory} = -7.61, SE = 0.59, p < .001$), and a significant Similarity*Condition interaction suggested that participants in the Memory condition believed that the more similar a team’s building style, the more similar a team’s building style, the more they could help Max remember the crucial trick ($\beta_{Sim.Mem} = 1.25, SE = 0.09, p < .001$).

General Discussion

Though it’s well-known that people will rewire their social networks in response to feedback — e.g., they “unfollow” inaccurate informants or uncooperative partners in favor of more reliable or cooperative connections (Almaatouq et al., 2020; Rand, Arbesman, & Christakis, 2011) — network-based approaches to social learning frequently abstract away from individual learners’ ability to control the structure of their networks. These abstractions often appear to tacitly assume that people are unwitting victims of the influence their networks can have on them. Yet, many of the psychological mechanisms driving network effects are highly intuitive: learning from one’s community, “echo chambers”, drift, recombining diverse ideas, to name a few. To the extent that people’s strategies for learning from others are guided by their intuitive theories about their utility (Heyes, 2019), intuitive theories about how networks shape the flow of information could allow people to control that flow by deliberately rewiring their networks. These decisions may not only impact their individual success as social learners, but could exert pressure on the evolution of the networks themselves. For instance, when selection favors skill specialists over skill generalists, networks grow denser over time — in other words, guilds form (Smolla & Akçay, 2019). It’s worth noting the congruity between this remarkable finding and a more familiar understanding of how guilds are formed: specialists form guilds deliberately in order to guarantee the quality of their work and training, as well as to protect trade secrets, just as they abandon guilds that become obstacles to innovation because cloistering has cut the flow of fresh ideas to a trickle (Aral & Alstynne, 2011). Thus, we suggest that research on people’s intuitive theories of network effects could further our understanding of how networks have shaped cumulative culture. Our results suggest that by adulthood, at least some of the basic mechanisms of network effects are indeed intuitively understood.

Acknowledgments

This research was funded by NSF grant DRL 1561143 to Frank Keil.

References

- Aral, S., & Alstynne, M. V. (2011). The Diversity-Bandwidth Trade-off. *American Journal of Sociology*, *117*(1), 82.
- Aboody, R., Zhou, C., & Jara-Ettinger, J. (2021). In Pursuit of Knowledge: Preschoolers Expect Agents to Weigh Information Gain and Information Cost When Deciding Whether to Explore. *Child Development*, *92*(5), 1919–1931.
- Baker, C. L., Jara-Ettinger, J., Saxe, R., & Tenenbaum, J. B. (2017). Rational quantitative attribution of beliefs, desires and percepts in human mentalizing. *Nature Human Behaviour*, *1*(4).
- Cappella, E., Neal, J. W., & Sahu, N. (2012). Children's Agreement on Classroom Social Networks: Cross-Level Predictors in Urban Elementary Schools. *Merrill-Palmer Quarterly*, *58*(3), 285–313.
- Derex, 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.
- Derex, M., Bonnefon, J.-F., Boyd, R., & Mesoudi, A. (2019). Causal understanding is not necessary for the improvement of culturally evolving technology. *Nature Human Behaviour*.
- Derex, M., & Boyd, R. (2016). Partial connectivity increases cultural accumulation within groups. *Proceedings of the National Academy of Sciences*, *113*(11), 2982–2987.
- Derex, M., & Mesoudi, A. (2020). Cumulative Cultural Evolution within Evolving Population Structures. *Trends in Cognitive Sciences*.
- Gest, S. D., Farmer, T. W., Cairns, B. D., & Xie, H. (2003). Identifying Children's Peer Social Networks in School Classrooms: Links Between Peer Reports and Observed Interactions. *Social Development*, *12*(4), 513–529.
- Gopnik, A., Frankenhuis, W. E., & Tomasello, M. (2020). Introduction to special issue: 'Life history and learning: how childhood, caregiving and old age shape cognition and culture in humans and other animals.' *Philosophical Transactions of the Royal Society B: Biological Sciences*, *375*(1803), 20190489.
- Hardwig, J. (1991). The Role of Trust in Knowledge. *The Journal of Philosophy*, *88*(12), 693–708.
- Harris, P. L., Koenig, M. A., Corriveau, K. H., & Jaswal, V. K. (2018). Cognitive Foundations of Learning from Testimony. *Annual Review of Psychology*, *69*(1), 251–273.
- Heyes, C. (2016). Who Knows? Metacognitive Social Learning Strategies. *Trends in Cognitive Sciences*, *20*(3), 204–213.
- Kempe, M., & Mesoudi, A. (2014). An experimental demonstration of the effect of group size on cultural accumulation. *Evolution and Human Behavior*, *35*(4), 285–290.
- Liberman, Z., Gerdin, E., Kinzler, K. D., & Shaw, A. (2020). (Un)common knowledge: Children use social relationships to determine who knows what. *Developmental Science*.
- Liu, L., Dehmamy, N., Chown, J., Giles, C. L., & Wang, D. (2021). Understanding the onset of hot streaks across artistic, cultural, and scientific careers. *Nature Communications*, *12*(1), 5392.
- Lutz, D. J., & Keil, F. C. (2002). Early understanding of the division of cognitive labor. *Child Development*, *73*(4), 1073–1084.
- Mason, W. A., Jones, A., & Goldstone, R. L. (2008). Propagation of innovations in networked groups. *Journal of Experimental Psychology: General*, *137*(3), 422–433.
- Mesoudi, A., & Thornton, A. (2018). What is cumulative cultural evolution? *Proceedings of the Royal Society B: Biological Sciences*, *285*(1880), 20180712.
- Smolla, M., & Akçay, E. (2019). Cultural selection shapes network structure. *Science Advances*, *11*.
- Zeng, A., Fan, Y., Di, Z., Wang, Y., & Havlin, S. (2021). Fresh teams are associated with original and multidisciplinary research. *Nature Human Behaviour*, *5*(10), 1314–1322.