



Research

Cite this article: Derex M, Bonnefon J-F, Boyd R, McElreath R, Mesoudi A. 2025 Social learning preserves both useful and useless theories by canalizing learners' exploration. *Proc. R. Soc. B* **292**: 20242499.

<https://doi.org/10.1098/rspb.2024.2499>

Received: 17 October 2024

Accepted: 19 December 2024

Subject Category:

Behaviour

Subject Areas:

behaviour, cognition

Keywords:

cultural evolution, cumulative culture, innovation, cultural inertia, social learning

Author for correspondence:

Maxime Derex

e-mail: maxime.derex@iast.fr

Electronic supplementary material is available online at <https://doi.org/10.6084/m9.figshare.c.7632413>.

Social learning preserves both useful and useless theories by canalizing learners' exploration

Maxime Derex^{1,2}, Jean-François Bonnefon^{1,2}, Robert Boyd^{3,4}, Richard McElreath⁵ and Alex Mesoudi⁶

¹Toulouse School of Economics, CNRS, and ²Institute for Advanced Study in Toulouse, University of Toulouse Capitole, Toulouse, France

³Institute of Human Origins, and ⁴School of Human Evolution and Social Change, Arizona State University, Tempe, AZ 85287, USA

⁵Max Planck Institute for Evolutionary Anthropology, Leipzig, Germany

⁶Human Behaviour and Cultural Evolution Group, Centre for Ecology and Conservation, University of Exeter, Penryn TR10 9FE, UK

MD, 0000-0002-1512-6496; AM, 0000-0002-7740-1625

In many domains, learning from others is crucial for leveraging cumulative cultural knowledge, which encapsulates the efforts of successive generations of innovators. However, anecdotal and experimental evidence suggests that reliance on social information can reduce the exploration of the problem space. Here, we experimentally investigate the extent to which cultural transmission fosters the persistence of arbitrary solutions in a context where participants are incentivized to improve a physical system across multiple trials. Participants were exposed to various theories about the system, ranging from accurate to misleading. Our findings indicate that even under conditions conducive to exploration, the transmission of cultural knowledge canalizes learners' focus, limiting their consideration of alternative solutions. This effect was observed in both the theories produced and the solutions attempted by participants, irrespective of the accuracy of the provided theories. These results challenge the notion that arbitrary solutions persist only when they are efficient or intuitive and underscore the significant role of cultural transmission in shaping human knowledge and technologies.

1. Introduction

In many domains, learning from others can provide valuable information about which solutions are worth considering and which are not [1–8]. This is especially true in the domain of technology. Technologies are typically the product of decades, centuries or even millennia of cumulative cultural evolution [9,10]. The technical solutions that surround us today embody the efforts of successive generations of innovators, and disregarding this accumulated knowledge to rely solely on our intuitions can have detrimental consequences [11,12].

Anecdotal evidence, however, suggests that learning from others can impede the discovery of better alternatives. For instance, a less-than-optimal turbine blade design in the early days of aircraft gas turbines went unnoticed for many years [13, p.187]. Turbine blades are heated by high-temperature exhaust gases resulting from the combustion of fuel in the engine. An early arbitrary decision was to position the blades' fixation point near the inner end of the airfoil. This caused the turbine disc, the part to which blades are attached, to overheat due to heat transfer. Consequently, engineers were compelled to use a type of steel that was dense, expensive and less reliable




Figure 1. Illustration of the physical system used in the experiment. The wheel had four radial spokes, and one weight could be moved along each spoke. The time it takes for the wheel to cover the track is determined by its Inertia and CoM. (a) A balanced wheel, consistent with the received Misleading theory, that does not properly exploit either Inertia or CoM. Here, the wheel has its centre of mass on the axis of rotation, and the vertical weights are farther away from the axis than the horizontal weights. (b) A balanced wheel, consistent with the Inertia theory, that solely exploits Inertia. B covers the track faster than A because of its lower Inertia. (c) An unbalanced wheel, consistent with the CoM theory, that exploits CoM. Here, the wheel does not have its centre of mass on the axis of rotation. C and A have comparable Inertia but C benefits from better acceleration because of its CoM. (d) A compact and unbalanced wheel, consistent with the Correct theory, that properly exploits both CoM and Inertia. D benefits from better acceleration than A and B because of its CoM, and a faster top speed than A and C because of its lower Inertia. Under the conditions of our experiment, D covers the track faster than A, B and C.

than alternatives. It took almost a decade to resolve this issue, by simply increasing the distance between the inner end of the airfoil and the blades' fixation point, thereby reducing heat transfer to the turbine disc [13].

Alongside anecdotal evidence, several experimental studies conducted among Western participants have revealed that learning from others might be detrimental [14–16]. For instance, fixation researchers, who study how new ideas originate, have shown that individuals inadvertently restrict the range of ideas they consider after being shown pictures of existing solutions [16]. Similarly, cognitive scientists have demonstrated that children who are told the function of a toy engage in more limited exploration and are less likely to discover alternative functions than children who are not told about the toy's function [14].

These results suggest that the transmission of cultural information might promote the persistence of arbitrary or suboptimal solutions by preventing social learners from thoroughly exploring the solution space. However, current research is limited in several critical ways. First, studies often involve scenarios where learners lack objective feedback about their performance and, crucially, are unable to iteratively refine their solutions. This neglects the potential impact of repeated trials and objective feedback, which could help learners discover more rewarding yet initially overlooked areas of the solution space, encouraging them to move beyond existing solutions [15]. Second, existing studies are typically constrained by tasks that do not allow for a detailed mapping of how social learners explore the solution space. Consequently, we do not know whether social information diminishes learners' overall exploration or merely channels it towards specific areas of the solution space.

Here, we aim to investigate whether, and if so how, cultural transmission promotes the persistence of arbitrary solutions in a context where participants are incentivized to improve a physical system across several trials. This requires our experimental task to exhibit two specific features. First, the task must provide participants with accurate and objective feedback on the performance of their solutions. Second, the task must be associated with a well-defined solution space, allowing us to analyse the effect of social information on participants' exploration patterns.

Our experimental task comprises a physical system with a wheel that travels down a 1 m long inclined track, as previously used in Derex *et al.* [17]. The wheel has four radial spokes each of which has a weight that can be moved along its length to one of 12 positions, creating a space of 20 736 unique configurations (figure 1). The aim for the participant is to position the spoke weights to minimize the time taken for the wheel to descend the track (video recordings available at <https://osf.io/g74y5>). The time it takes for the wheel to do this is determined by two variables: its moment of inertia (henceforth 'Inertia') and the position of its centre of mass ('CoM'). The inertia of the wheel depends on how mass is distributed around the axis. The wheel has lower Inertia and will rotate more easily when weights are closer to the axis of rotation. Asymmetrical wheels do not have their CoM on the axis of rotation, which can give wheels better initial acceleration. These two variables imply that four qualitatively distinct types of theories can be generated about the physical system: CoM theories consider CoM but ignore Inertia (e.g. 'The wheel covers the track faster when its top and right weights are farther from the axis than its bottom and left weights'); Inertia theories consider Inertia but ignore CoM (e.g. 'The wheel covers the track faster when all its weights are close to the axis'); Correct theories consider both Inertia and CoM (e.g. 'The wheel covers the track faster when all its weights, except the top one, are close to the axis') and Misleading theories consider neither CoM nor Inertia (e.g. 'The wheel covers the track faster when its vertical weights are farther away from the axis than its horizontal weights').

Participants ($n = 200$) were exposed to a range of theories about our physical system that were generated by participants exposed to the same physical apparatus in a previous experiment [17]. Participants were randomly assigned to either a Control treatment or one of the four Social Information treatments (40 participants per treatment). Participants assigned to a Social Information treatment received a theory that was either Correct, Inertia-related, CoM-related or Misleading (see Methods for details). Control participants were exposed to no theory. This allowed us to study the persistence of existing theories and their effects on social learners' exploration patterns.

The experiment was preregistered (<https://osf.io/y2xaz>) and organized as follows: all participants were first asked to choose a configuration and then asked to write a theory about what makes the wheel cover the track in the shortest amount of time.




Figure 2. Overview of the experimental procedure. Naïve period: Participants rely solely on their prior knowledge. They choose a configuration (t1) and then write a theory about what makes the wheel cover the track in the shortest time (t2). Social information period: All participants, except those in the control treatment, receive one of the four types of theories formulated by a participant who was exposed to the same task in a previous experiment. Exploration period: Participants interact with the physical system for five trials. At each trial, they choose a configuration (t1), the wheel is released, and participants are automatically provided with their wheel's average speed and the associated pay-off (t2). Test period: Participants choose a bonus configuration, whose associated pay-off is multiplied by 3 (t1) and then write a potentially updated theory about what makes the wheel cover the track in the shortest time (t2). Finally, the wheel is released, and participants are provided with their wheel's average speed and the associated pay-off (t3).

During this naïve period, participants relied solely on their prior knowledge or intuition, as they had not yet observed the wheel going down the track (figure 2). Then, participants from the social information treatments received one of the four types of theory. All participants were then given the opportunity to change the configuration they had initially chosen. This was followed by an exploration period during which participants had five successive trials to optimize their wheel and maximize their pay-off. After each trial, participants were automatically provided with their wheel's average speed and the associated pay-off (range: €0–3). After completing five trials, participants moved to the test period. They were invited to choose a bonus configuration whose associated pay-off was multiplied by three (range: €0–9) before being asked again to provide a theory about the wheel. Finally, they were provided with their final wheel's average speed and the associated pay-off.

Our main preregistered hypotheses were that: (i) social learning promotes the persistence of whichever theory is received because (ii) those received theories canalize learners' exploration and prevent them from thoroughly exploring the solution space.

2. Results

Theories produced during the naïve and test periods were categorized by human raters according to whether they harness the effects of Inertia and/or the CoM of the wheel. Theories that considered only Inertia were categorized as 'Inertia'. Theories that considered only CoM were categorized as 'CoM'. Theories that considered both were categorized as 'Correct'. Theories that overlooked both relevant variables were categorized as 'Misleading' if they were incorrect in a manner consistent with the received Misleading theory, otherwise as 'Others'.

(a) Understanding patterns among naïve participants

Among the different types of theories formulated by naïve participants from the Control treatment, 9/40 (0.225) were categorized as CoM and 4/40 (0.1) were categorized as Inertia. No naïve participants from the control treatment were able to formulate a Correct theory at this stage of the experiment. 27/40 theories (0.675) overlooked both relevant variables. Among these, 1/40 (0.025) was categorized as Misleading and 26/40 (0.65) were categorized as Others. Of the 26 theories categorized as 'Others', 6 were incorrect regarding the effect of inertia (e.g. 'The wheel covers the track faster when all its weights are far from the axis'). The remaining theories were either insufficiently informative (e.g. 'The wheel covers the track faster when it is balanced/unbalanced'; 6 and 4 respectively) or unhelpful (e.g. 'The wheel covers the track faster when the weights propel it').

The probabilities of producing each type of theory were comparable between participants from the control treatment and participants from the social information treatments prior to receiving a theory (CoM: diff. in prob. 95% CI [-0.12, 0.16], mean = 0.03; Inertia: diff. in prob. 95% CI [-0.16, 0.04], mean = -0.04; Correct: diff. in prob. 95% CI [-0.02, 0.04], mean = 0.01; Misleading: diff. in prob. 95% CI [-0.08, 0.01], mean = -0.02; Others: diff. in prob. 95% CI [-0.13, 0.18], mean = 0.02).

(b) Effect of learning in the control treatment

After interacting with our physical system, the Inertia theory was most common among participants from the Control treatment (17/40 = 0.425), representing a reliable positive change compared with when participants were naive (difference in prob. 95% CI [0.15, 0.50], mean = 0.32). Theories categorized as Others were the second most common type produced by experienced participants (16/40 = 0.40), which represented a reliable decrease compared with when participants were naive (diff. in prob. 95% CI [-0.46, -0.03], mean = -0.25). The probabilities of producing the CoM, Correct and Misleading theories were not affected by interacting with the physical system (CoM: 7/40 = 0.175, diff. in prob. 95% CI [-0.23, 0.14], mean = -0.05; Correct: 0/40, diff. in prob. 95% CI [-0.02, 0.02], mean = 0; and Misleading: 0/40, diff. in prob. 95% CI [-0.09, 0.01], mean = -0.02).

(c) Effect of social information

Participants who received a theory were more likely to change their wheel's initial configuration than participants from the control treatment, indicating that participants were influenced by the social information they received (Control: 0.15; CoM: 0.40, diff. in prob. 95% CI [0.05, 0.43], mean = 0.25; Inertia: 0.45, diff. in prob. 95% CI [0.10, 0.48], mean = 0.29; Correct: 0.43, diff. in prob. 95% CI [0.08, 0.45], mean = 0.27; Misleading: 0.35, diff. in prob. 95% CI [0.01, 0.38], mean = 0.20). However, in accordance with our preregistered hypothesis, individuals' probability of changing their initial configuration was statistically comparable between social information treatments, indicating that the relevance of the received theories for reaching higher performances did not affect individuals' willingness to take social information into account (e.g. Misleading versus Correct: diff. in prob. 95% CI [-0.13, 0.28], mean = 0.07). Contrary to one of our secondary preregistered hypotheses, individuals' probability of changing their initial configuration was not lower in men compared with women (Men: 0.43; Women: 0.39; diff. in prob. 95% CI [-0.18, 0.12], mean = -0.04).

We now look at the effect of interacting with the task among participants who received social information. Our results confirm our preregistered hypothesis that receiving a theory about how a task works before interacting with the task increases the likelihood that individuals will produce the same theory after interacting with the task, compared with participants who do not receive a theory (figure 3). The probability of producing the Correct theory after interacting with the task was 0 among participants from the Control treatment and 0.45 among participants who received the correct theory, which represents a reliable increase compared with participants from the Control treatment (diff. in prob. 95% CI [0.29, 0.59], mean = 0.44). Reliable increases in the probability of producing the theory received were also observed when participants received the Inertia theory (from 0.43 to 0.63, diff. in prob. 95% CI [-0.01, 0.41], mean = 0.20), the CoM theory (from 0.18 to 0.45, diff. in prob. 95% CI [0.07, 0.45], mean = 0.27) and the Misleading theory (from 0 to 0.30, diff. in prob. 95% CI [0.17, 0.44], mean = 0.30).

Qualitatively similar results are obtained when we infer participants' theories from their final wheel configuration (figure 4). A reliable increase in the probability of producing a wheel consistent with the received theory is observed when participants received the Correct theory (from 0 to 0.47, diff. in prob. 95% CI [0.31, 0.62], mean = 0.44), the Inertia theory (from 0.30 to 0.48, diff. in prob. 95% CI [-0.03, 0.38], mean = 0.17), the CoM theory (from 0.28 to 0.65, diff. in prob. 95% CI [0.16, 0.56], mean = 0.37) and the Misleading theory (from 0.13 to 0.40, diff. in prob. 95% CI [0.09, 0.44], mean = 0.27).

Among participants in the Correct treatment, the probability of writing a theory categorized as Correct was comparable to the probability of producing a configuration consistent with the Correct theory (0.45 and 0.47, respectively, diff. in prob. 95% CI [-0.23, 0.18], mean = -0.03). This contradicts our secondary preregistered hypothesis that correct, two-dimensional theories might fail to be properly transmitted through text messages because participants would tend to write simpler, unidimensional theories that are less costly to articulate.

(d) Effect of social information on exploration patterns

Over the five trials of the exploration period, participants in the Control treatment produced an estimated mean of 4.75 unique configurations (95% CI [4.57, 4.94]), which is comparable to the levels of exploration exhibited by participants who received a theory (contrasts with: Misleading: 95% CI [-0.37, 0.12], mean = -0.13; CoM: 95% CI [-0.20, 0.30], mean = 0.05; Inertia: 95% CI [-0.35, 0.15], mean = -0.10; Correct: 95% CI [-0.20, 0.30], mean = 0.05).

However, consistent with our preregistered hypothesis, exploration patterns were influenced by the type of theory received (figure 5). Participants who received the Misleading theory mostly produced balanced wheels with their horizontal weights closer to the axis than vertical weights (figure 5a,e). Participants who received the CoM theory mostly produced unbalanced wheels with their top and right weights at the outermost position (figure 5b,f). Participants who received the Inertia theory mostly produced balanced wheels with all their four weights close to the axis (figure 5c,g). Participants who received the Correct theory mostly produced unbalanced wheels with all their weights, except the top one, close to the axis (figure 5d,h).

Exploratory analyses indicate that participants in the Misleading and Inertia treatments reliably produced more balanced wheels (Misleading: 95% CI [2.68, 3.50], mean = 3.09; Inertia: 95% CI [2.64, 3.45], mean = 3.05) compared with unbalanced wheels (Misleading: 95% CI [1.18, 2.02], mean = 1.60; Inertia: 95% CI [1.16, 2.04], mean = 1.60). This suggests a greater focus on varying the moment of inertia of the wheel rather than the position of its centre of mass. In the CoM and Correct treatments, the opposite exploration pattern was observed (figure 6). Participants in these treatments reliably produced fewer balanced wheels (CoM: 95% CI [0.55, 1.37], mean = 0.95; Correct: 95% CI [0.74, 1.56], mean = 1.16) compared with unbalanced wheels (CoM: 95%

Received theories are sticky

Participants who received a theory are more likely to produce **this theory** than participants from the control group. Participants' theories were categorized by coders based on their free-text responses.




Figure 3. Difference in the probability of producing each type of theory as a function of the theory received, compared with receiving no theory. Participants in the Control treatment (no theory) produced one out of five types of theories after interacting with our physical system: Correct (0%), Inertia (42.5%), CoM (17.5%), Misleading (0%) and Others (40%). The figure illustrates the difference in the probability of producing each type of final theory as a function of the theory received, compared with those values. For instance, the left column illustrates the difference in the probability of producing each type of theory between participants in the Control treatment and participants who received the Correct theory. When participants received the Correct theory, the probability of producing each type of theory changed as follows: Correct went from 0 to 0.45 (+0.45), Inertia from 0.42 to 0.20 (−0.22), CoM from 0.17 to 0.20 (+0.03), and Misleading from 0 to 0 (0), respectively. Red dots along the diagonal indicate that, compared with participants in the Control treatment, receiving any theory before interacting with the task increases individuals' probability of producing the same theory after completing the task.

Final wheel configurations bear the footprint of received theories

Analyses based on participants' final wheel configuration show qualitatively similar results compared to coders categorizations (see Figure 3). Here participants' theories are automatically inferred from their wheel configurations, instead of having coders categorizing participants' free-text responses.




Figure 4. Difference in the probability of producing a configuration consistent with each type of theory as a function of the theory received, compared with receiving no theory. Red dots along the diagonal indicate that, compared with participants in the Control treatment, receiving any theory before interacting with the task increases individuals' probability of producing a final wheel configuration consistent with the theory received.

CI [3.41, 4.29], mean = 3.85; Correct: 95% CI [3.21, 4.08], mean = 3.65), indicating a greater focus on varying the position of the wheel's centre of mass rather than its moment of inertia.

(e) Effect of social information on performance

Participants in the Inertia and Correct treatments performed better in the bonus trial than those in the Control treatment (Inertia 95% CI [0.00, 1.86], mean = 0.93; Correct 95% CI [0.00, 1.85], mean = 0.92). In contrast, participants in the CoM and Misleading treatments performed comparably to the Control group (CoM 95% CI [−1.55, 0.28], mean = −0.63; Misleading 95% CI [−1.16, 0.70], mean = −0.22). This partially contradicts our preregistered hypothesis that partially and fully correct theories would be beneficial, while the misleading theory would be detrimental.

Expanding the analysis to all trials reveals that participants improved their performance across trials in all treatments (Misleading : 95% CI [0.09, 0.17], mean = 0.13; CoM: 95% CI [0.15, 0.20], mean = 0.17; Inertia: 95% CI [0.15, 0.21], mean = 0.18; Correct: 95% CI [0.15, 0.22], mean = 0.18). Across all trials, the performance of participants who received a theory differed from that of participants in the Control treatment, with the exception of the CoM treatment (contrast with Control treatment 95% CI

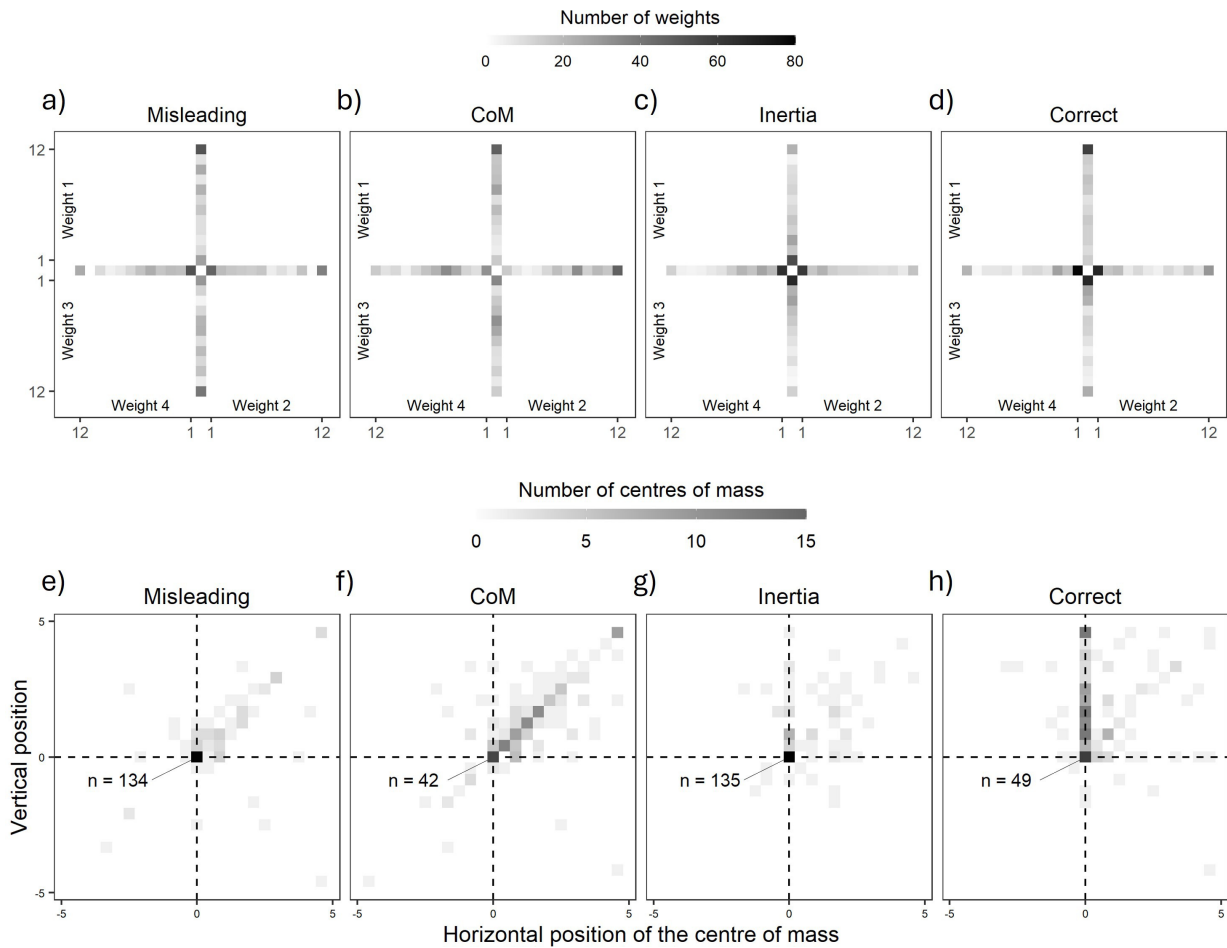


Figure 5. Exploration patterns are affected by the type of theory received. (a–d) Heat maps illustrating the most frequent weight positions along each spoke in each treatment. (e–h) Most frequent positions of the wheels' centre of mass in each treatment. Wheels with both their horizontal and vertical weights equidistant from the axis are balanced, with their centre of mass located on the axis of rotation (i.e. at the intersection of the dotted lines in panels (e) to (h)). Values 1–12 in panels (a) to (d) describe the positions of weights 1–4. Values –5 to +5 in the bottom panel describe the x - and y -coordinates of the wheels' centre of mass.

Receiving a theory analyses exploration

High levels of exploration are observed among social learners but exploration patterns are affected by the type of theory received. Participants from the Misleading and Inertia treatments produced more *balanced wheels* than *unbalanced wheels*, while the opposite pattern is observed among participants from the CoM and Correct treatments




Figure 6. Receiving a theory affects the exploration of dimensions relevant to the performance of the wheel. Errors bars show 95% CIs. The horizontal line indicates the maximum number of unique configurations that can be produced during the exploration period (see figure 1 for examples of balanced and unbalanced wheels).

[–0.17, 0.16], mean = –0.01; Trial \times Treatment 95% CI [–0.05, 0.02], mean = –0.01). Both the Inertia and Correct treatments resulted in better performance (Inertia: contrast with Control treatment 95% CI [0.05, 0.27], mean = 0.16; Trial \times Treatment 95% CI [–0.03, 0.01], mean = –0.01; Correct: contrast with Control treatment 95% CI [0.01, 0.18], mean = 0.10; Trial \times Treatment 95% CI [–0.02, 0.01], mean = 0.00). Participants who received the Misleading theory displayed slower improvement than participants from the Control treatment (contrast with Control treatment 95% CI [–0.14, 0.49], mean = 0.18; Trial \times Treatment 95% CI [–0.12, 0.01],

mean = -0.05). Additional analyses indicate that the relatively poor performance of participants in the CoM treatment results from a reliably higher probability of producing wheels that did not complete a full first revolution due to extreme positions of the wheel's centre of mass (electronic supplementary material, table S1).

3. Discussion

Here we experimentally investigated whether cultural transmission promotes the persistence of arbitrary solutions by preventing social learners from thoroughly exploring the solution space. Specifically, we tested whether exposing participants to different theories about a physical system—a wheel descending a track—affected their exploration patterns in a way that promotes the persistence of these theories. Our experiment reveals that, despite participants being incentivized to produce accurate theories, there was a reliable increase in the likelihood of producing the theory that was received from a previous participant. The impact of social information was evident not only in the theories produced but also in the wheel configurations that were generated by participants.

These results support recent experimental findings highlighting what has been termed the double-edged sword of pedagogy, where teaching increases the likelihood of performing relevant behaviours but also reduces the likelihood of discovering alternative ones [14]. However, our study offers a different perspective on how social information affects learners' exploration. While previous research has suggested that social learning limits exploration, our results indicate that social learning canalizes learners' exploration without necessarily reducing its overall extent. The tasks used in earlier experiments may have masked similar effects. For instance, in Bonawitz *et al.*'s [14] study, children were given unfamiliar toys with four non-obvious functions, such as squeaking, flashing lights and a hidden mirror. Children who were informed about one function of the toy were less likely to discover the other functions compared with those who were not given any information [14]. However, the study also found that children who were told about a function spent more time using that demonstrated function. This increased time spent was interpreted as evidence of limited exploration. This interpretation assumes that discovering only the different functions was the sole objective. Yet, it seems reasonable to assume that by spending more time using the demonstrated function, children might have discovered additional properties of that function (for instance, how to modulate the intensity or tone of the squeaking function).

Our results also indicate that the effect of social information on learners' exploration extends beyond situations where learners are provided with useful information. Existing evidence of the detrimental effect of social information predominantly comes from conditions where social learners were given viable solutions (but see [15]). Consequently, previous work could not rule out the possibility that participants assumed there was nothing more to discover [14]. In our study, participants were provided with one of the four qualitatively distinct types of theory, ranging from fully correct to misleading. Moreover, participants were incentivized to improve their solutions, had multiple opportunities to explore alternative solutions and could immediately and accurately assess the efficiency of their solutions. Despite this, we observed reliable increases in the probability of producing both a theory and configurations consistent with the theory received, even when the theory was only partially correct or misleading.

These results are particularly surprising given that participants exhibited high levels of exploration across all treatments. Under these conditions, one might expect that extensive exploration combined with non-noisy pay-offs would enable participants to quickly dismiss inaccurate theories. For instance, in an experiment based on a simple multi-armed bandit task where the usefulness of social information was varied, data show that participants tended to disregard the received solution if its pay-off proved worse than the solutions participants discovered themselves through exploration [15]. However, in our experiment, social information channelled learners' exploration in a way that made it difficult for them to diverge from the theory they received. Indeed, receiving a theory led learners to produce configurations that were mostly consistent with the theory received. This tendency made individuals less likely to generate configurations that could challenge the theory they received thus reducing the likelihood of discovering the effects of variables not emphasized by the received theory.

One potential mechanism by which cultural transmission might influence the range of solutions individuals explore is by shaping their representation of the problem. This phenomenon, often highlighted in literature on cognitive flexibility, suggests that individuals may struggle to shift their focus from specific features of a problem to alternative ones [18–20]. For example, experiments have shown that when individuals are presented with a series of single digits and tasked with either determining if the digit is odd or even, or if it is larger or smaller than five, they respond more quickly when repeating the same task than when switching tasks [21]. In more complex conditions, the specific features participants focus on guide their actions, potentially confining them to a subset of the search space where the most optimal solution cannot be found [18,20]. In these experiments, the representations held by learners typically arise from their direct interaction with the task. However, our findings suggest that the cultural transmission of information can produce a similar effect. This may help explain why social learners improved their solutions across all treatments. Received theories direct attention to specific features (such as the relative position of vertical weights versus horizontal weights in our Misleading treatment), thus narrowing the set of possible solutions individuals consider when improving their wheel. While this limits the search to certain areas of the solution space, it still allows participants to discover more efficient solutions.

Our results contribute to the debate on humans' propensity to rely on social information and the implications for the persistence of arbitrary solutions. Some scholars argue that individuals are inclined to rely on social information, which can facilitate the adoption of hard-to-devise, unintuitive solutions [12,22,23]. Others contend that individuals should be cautious about social information to avoid the risk of being accidentally or intentionally misled [24,25]. According to the former view, arbitrary cultural solutions are likely to persist, while the latter suggests these solutions should quickly fade unless they are

particularly intuitive to individuals. Our experiment supports the former view, challenging the idea that arbitrary solutions persist only when they are intuitive to participants. Notably, only 1 out of 40 participants produced a theory categorized as Misleading during the naive period, and none produced a theory categorized as Correct. Despite this, both types of theory persisted in our experiment.

One might argue that our experimental design was biased towards observing a persistent impact of social information. Indeed, theory predicts that individuals are likely to heavily rely on social information when they are uncertain [26], either because they have no relevant prior information [1], because the number of potential solutions is large [27], or because others possess more reliable information [28]. In our experiment, participants faced an unfamiliar task with a large solution space and were informed that the theories they received came from individuals with prior experience with the physical system. However, we believe these conditions appropriately reflect the challenges individuals face when attempting to improve existing technologies. Moreover, at least two aspects of our design make the experiment conservative. First, our task was low-dimensional (i.e. only Inertia and CoM affect the dynamics of the wheel), which is likely to reduce uncertainty, and thereby reliance on social information, compared with facing a real technology, which tends to be high dimensional. Second, we used non-noisy pay-offs, whereas most real technologies provide noisy pay-offs, known to increase the persistence of arbitrary solutions [29].

In conclusion, our experiment demonstrates that the transmission of cultural knowledge can act as a cognitive barrier, hindering individuals from thoroughly exploring the solution space. This finding aligns with the broader framework of cultural evolution, which emphasizes the role of cultural transmission in shaping human behaviour across domains [11,30–32]. More broadly, our results highlight the complex interplay between cultural transmission and individual and collective exploration. Understanding this dynamic is crucial not only for understanding patterns of cultural evolution but also for designing effective strategies and interventions that enable us to reap the collective benefits of social learning (i.e. cumulative culture and collective intelligence) while mitigating the associated costs (i.e. fixation effects and canalized exploration). While we demonstrated that receiving written theories can promote the persistence of arbitrary solutions, it is important to note that the effects observed in our study emerged under specific conditions: Western participants received written theories and solved the problem in isolation. These effects may differ or be mitigated when alternative forms of social information are used, such as observation, demonstration or direct interaction, or within different group or network structures. Group connectivity patterns, for example, are known to positively influence cultural evolution by promoting exploration, boosting creativity and facilitating the recombination of solutions [33–35]. Additionally, cultural variability may shape how social information influences exploration, highlighting the need for further investigation into the interplay between cultural transmission and problem-solving across diverse contexts. Future work should aim to better understand the potential constraints imposed by cultural transmission and explore approaches to promote a more balanced and diverse exploration of solutions.

4. Methods

(a) Participants

In total, 200 participants took part in the study (100 women and 100 men). Participants were randomly selected from a database managed by Catholic University of Lille and recruited by email from various universities in Lille, France. The participants ranged in age from 18 to 50 years (mean of 21.2, s.d. of 3.96). Participants received €3 for participating and an additional amount ranging from €0 to €25 depending on their performance (see below).

(b) Ethical statement

The study was carried out in accordance with the ethical standards of the 1964 Declaration of Helsinki and the guidelines of the British Psychological Society's Code of Human Research Ethics. All methods were approved by the University of Exeter Biosciences Research Ethics Committee (2019/1940). All participants provided written, informed consent before taking part in the experiment.

(c) Experimental apparatus

The experimental apparatus was similar to that used in [17]. It consisted of a wheel that had to travel down a 1 m long inclined track. The wheel had four radial spokes, and one weight could be moved along each spoke. Weights could be placed on one of the 12 discrete positions which created a space of 20 736 unique configurations. The performance of the wheel depends on two variables: its moment of inertia and the position of its centre of mass. The wheel's moment of inertia depends on how mass is distributed around its axis of rotation. Wheels with a smaller moment of inertia (i.e. wheels that have their weights closer to the axis) require less torque to increase angular momentum and spin faster (Video recordings are available at <https://osf.io/g74y5/>). Asymmetrical wheels do not have their centre of mass on the axis of rotation, which can provide a better initial acceleration. When the centre of mass of the wheel is in the wheel's upper right quadrant (assuming the wheel goes downhill from left to right), more potential energy is converted into angular kinetic energy so that the wheel will benefit from higher increases in angular momentum. In our experiment, both the wheel's moment of inertia and the position of its centre of mass had to be taken into account to reach the best performance. A higher centre of mass can produce better acceleration, but it will increase the wheel's moment of inertia and so there was a trade-off between maximizing acceleration and minimizing inertia [17].

(d) Procedure

The experiment took place in an experimental room at the Laboratory for Experimental Anthropology at Catholic University of Lille. For each approximately 20 min session, a single individual was recruited and sat at a computer that was placed parallel to and at 2 m from the experimental apparatus. Participants were randomly assigned to one condition of the experiment. Before starting the experiment, participants were asked to sign a consent form and were asked their age. At the end of the experiment, participants received a reward according to their performance. Participants entered and left the room by two different doors to prevent any form of direct interactions between participants.

(e) Experimental design

The experiment comprised three distinct periods: a naive period, an exploration period and a test period. Participants chose their configurations through a computer program using four sliders.

(i) naive period (Trial 0)

Participants started by choosing a configuration. Right after participants confirmed their configuration, they were asked to write their theory about what makes the wheel covering the distance in the shortest amount of time. Theories had to be less than 340 characters long and always started with 'The wheel covers the distance faster when...' in order to encourage participants to provide a general statement about the wheel. During this period, participants relied solely on their prior knowledge as they had not yet observed the wheel being released. Participants then received one of the five experimental treatments before being given the opportunity to change the configuration they had initially chosen.

(ii) Exploration period (Trials 1–5)

Once participants confirmed the configuration of their wheel (whether they had changed it or not), the experimenter positioned the weights on the physical wheel accordingly (the computer screen was projected onto a wall to the right of the participant in order to allow the experimenter to see the chosen configuration without interacting with the participant). The wheel was then positioned on the rails. A mechanical lever maintained the wheel motionless, with two of its spokes parallel to the ground at its starting position. Once released, the time it took the wheel to descend the track was automatically recorded by the computer program, and the wheel's average speed and associated pay-off were automatically displayed on the participant's screen. The participant could then choose a new configuration. The procedure was repeated until participants had completed their five trials.

(iii) Test period (bonus configuration)

After having completed their five trials, participants were invited to choose a bonus configuration whose associated score is multiplied by 3. After confirming their configuration, participants were asked to provide a theory about the wheel. Only after providing their theory, participants could observe the wheel going down.

(iv) Experimental treatments

Five treatments were run. All participants except those in the control treatment were provided with social information. Social information took the form of a theory that was produced by a participant from a previous experiment involving the same physical apparatus. Four theories that vary in their accuracy were chosen in order to cover the three qualitatively distinct types of theories that participants can receive: fully correct, partially correct and misleading. In the Correct treatment, participants were provided with a theory that states: 'The wheel covers the distance faster when all weights are close to the axis except the top weight that has to be a bit farther away'. This theory encourages participants to produce wheels with a low moment of inertia and a centre of mass located above the wheel's axis of rotation (at the wheel's initial position). This theory captures the two principles that allow participants to produce the most efficient wheels. In the Inertia treatment, participants were provided with a partially correct theory that emphasizes solely the role of the wheel's moment of inertia: 'The wheel covers the distance faster when the weights are close to the axis'. In the Centre of Mass treatment, participants were provided with another partially correct theory, the one that emphasizes solely the role of the position of the wheel's CoM: 'The wheel covers the distance faster when the top and right weights are farther from the axis than the bottom and left weights'. In the Misleading treatment, participants were provided with a theory that does not emphasize either of the two dimensions that are relevant to the performance of the wheel: 'The wheel covers the distance faster when the horizontal weights are closer to the axis than the vertical weights'. This misleading theory was chosen because (i) it produces a recognizable wheel configuration that does not overlap with configurations consistent with the other seeded theories, ensuring we could identify the effect of social information and (ii) it was rare in the pre-existing

dataset, minimizing the likelihood that its ‘stickiness’ was due to participants’ prior intuitions rather than the influence of social information.

Theories were provided to participants at the end of the naive period. Participants were given the opportunity to change their initial configuration at the beginning of the exploration period in all treatments (including individuals from the control treatment who did not receive any theory). Theories were removed from participants’ screen after they validated their second configuration.

(v) Pre-experiment information

Instructions could be read on a computer screen and stated that the participants’ task was to position four weights on a wheel in order to minimize the time it takes the wheel to cover an inclined track. Participants were informed that they will have five trials to do this, and that their pay-off will be determined by the performance of each of their wheels. Additionally, they were told that they will have to formulate a theory about the wheel and that this theory will be evaluated and will determine part of their score. Participants were informed that they might be provided with a previous participant’s theory and that this theory might help them maximizing their score.

(vi) Participants’ pay-off

During the exploration period, the following equation determined the pay-off of each wheel:

$$[1 - ((\text{MaxSpeed} - \text{RecordedSpeed})/(\text{MaxSpeed} - \text{MinSpeed}))] \times 3 + \text{Bonus}$$

with MaxSpeed = 160 and MinSpeed = 96. RecordedSpeed was the recorded average speed of the wheel. Bonus took the value 0.2 for wheels that descended the rails and 0 otherwise. During the test period, wheels’ pay-off was multiplied by 3. Theories formulated by participants were immediately evaluated by the experimenter (but later independently coded for the purposes of statistical analyses, see below) and provided participants with an additional pay-off of €0, €0.5 or €1, depending on whether they mention none, one or two of the dimensions that are relevant to the performance of the wheel. Participants’ final pay-off corresponded to the sum of the pay-off of each of their wheels plus the pay-offs associated with their theories.

(vii) Theory coding (human raters)

For the purposes of statistical analyses, theories were coded by two sets of three individuals blind to the research question. Coders were explained the dynamics of the wheel (i.e. the respective role of the inertia and centre of mass in the performance of the wheel) before completing their task. The first set of coders were asked to code participants’ theories according to whether they contain accurate information related to the moment of inertia (Inertia) and/or centre of mass (CoM). A theory contained information related to the moment of inertia when it says that the wheels goes faster when its weights are close to the axis (e.g. ‘*The wheel covers the track faster when its weights are balanced and close to the axis*’). A theory contained information related to the centre of mass when it says that the wheel goes faster when its centre of mass is in the upper right quadrant (e.g. ‘*The wheel covers the track faster when its top and right weights are far from the axis and its bottom and left weights are close to it*’). Theories that contained accurate information about the effect of the Inertia of the wheel were considered as partially correct and consistent with the seeded Inertia theory. Theories that contained accurate information about the effect of the CoM of the wheel were considered as partially correct and consistent with the seeded CoM theory. Theories that contained accurate information about both effects were considered as correct and consistent with the seeded Correct theory. The second set of coders evaluated whether theories were consistent with the seeded Misleading theory. Theories that were not categorized as Correct, CoM, Inertia or Misleading were categorized as Others. A majority among coders determined final coding. Cohen’s kappa coefficients reveal either substantial or almost perfect agreement between raters (0.84 for Inertia, 0.78 for CoM and 0.96 for Misleading).

(viii) Theory coding (classification algorithm)

Participants’ configuration that was chosen just before formulating their updated theory (i.e. the bonus trial associated with the test period) was used to infer participants’ theories at the end of the experiment. The most compact configuration was considered as evidence for partially correct and ‘Inertia-related’ theory. Configurations with their centre of mass in the top right quadrant, minus those that are better than the most compact wheel, were considered as evidence for partially correct and ‘CoM-related’ theory. All configurations that had their bottom and left weight at the closest position to the axis and that were better than the most compact wheel were considered as evidence for accurate theory (electronic supplementary material, table S2). All configurations that had their horizontal weights closer to the axis than their vertical weights (minus those that correspond to either the CoM, Inertia or Accurate theories) were considered as evidence for Misleading theory. Other configurations were categorized as Others.

(f) Statistical analyses and models output

We ran a series of Bayesian models in R [36]. Models were fitted using the *rethinking* package [37] and 95% credible intervals were used to make inferences. Full details about our models are available at <https://doi.org/10.17605/OSF.IO/G74Y5>.

(i) Probability of changing configuration after being exposed to social information.

Configurations from the naive period and the first trial of the exploration period were used. We fitted a logistic regression with 'Change' as the response variable and one dummy variable for each social information treatment as predictor variables.

(ii) Probability of changing configuration among men and women

Configurations from the naive period and the first trial of the exploration period were used. We fitted a logistic regression with 'Change' as the response variable and one dummy variable for 'Women' as predictor variable.

(iii) Effect of individual learning (Control treatment) on the theory produced.

Theories from the naive and test periods were used. We fitted a categorical model where response ('Type of theory') and predictor ('naive') were all categorical and unordered. The model estimates the probability of producing each of the theories, which corresponds to a vector of probabilities, and lets this vector vary by treatment. The prior for the vector is a Dirichlet distribution. The Dirichlet distribution is a distribution for probabilities (with values between zero and one) that all sum to one. We used the same value for each variable of the vector, which corresponds to a uniform prior. To test for differences among categorical predictors, we calculated the differences between each contrast of interest and computed the highest posterior intervals from the distribution of these differences to make inferences.

(iv) Effects of social information on the theory produced

Theories from the test period were used. We fitted a categorical model similar to the one described above except that 'Treatments' was the predictor.

(v) Effect of social information on the number of unique configurations produced

Configurations from the exploration period were used. We fitted a linear regression with 'Number of unique configurations' as the response variable and one dummy variable for each treatment as predictor variables.

(vi) Effect of social information on the number of balanced and unbalanced configurations

Configurations from the exploration period were used. We fitted a linear regression with 'Number of unique balanced configurations' as the response variable and one dummy variable for each social information treatment as predictor variables. We fitted a similar model with 'Number of unique unbalanced configurations' as the response variable.

(vii) Effect of social information on performance

Configurations from the exploration and test periods were used. We fitted a linear model with 'Pay-off' as the outcome variable, 'Trial', 'Treatment' and 'Trial:Treatment' as predictor variables and 'Participant's identity' as random effect.

(viii) Deviation from pre-registered analyses

The pre-registered analysis on the effect of social information on performance uses 'Pay-off' as outcome variable instead of 'Speed'. 'Pay-off' and 'Speed' are linearly correlated, but using 'Pay-off' as the outcome variable allows the model to sample more efficiently and provides narrower confidence intervals (see electronic supplementary material in [17]). This is because the gap between the wheels that did not descend the rails and those that did is proportionally smaller when considering Pay-off (0 to €0.3–3) than Speed (0 to 98.2–154.7 m h⁻¹).

(g) Pre-registration

The study was pre-registered ([pre-registration 1](#); [pre-registration 2](#)).

Ethics. All methods were approved by the University of Exeter Biosciences Research Ethics Committee (2019/1940). All participants provided written, informed consent before taking part in the experiment.

Data accessibility. Codes used in this paper and data that support the findings of this study are available at [38]. Supplementary material is available online [39].

Declaration of AI use. We have not used AI-assisted technologies in creating this article.

Authors' contributions. M.D.: conceptualization, formal analysis, funding acquisition, investigation, methodology, project administration, writing—original draft, writing—review and editing; J.-F.B.: conceptualization, visualization, writing—review and editing; R.B.: conceptualization, writing—review and editing; R.M.: formal analysis, writing—review and editing; A.M.: conceptualization, supervision, writing—review and editing. All authors gave final approval for publication and agreed to be held accountable for the work performed therein.

Conflict of interest declaration. We declare we have no competing interests.

Funding. M.D. acknowledges Institute for Advanced Study in Toulouse (IAST) funding from the French National Research Agency (ANR) under grant no. ANR-17-EURE-0010 (Investissements d'Avenir programme) and ANR funding through the ANR-21-CE28-0019-01 OPTILEARN project.

Acknowledgements. We thank the editor and two anonymous reviewers for valuable comments on a previous version.

References

- Boyd R, Richerson PJ. 1985 *Culture and the evolutionary process*. Chicago, IL: University of Chicago Press.
- Boyd R, Richerson PJ, Henrich J. 2011 The cultural niche: why social learning is essential for human adaptation. *Proc. Natl Acad. Sci. USA* **108**, 10918–10925. (doi:10.1073/pnas.1100290108)
- Rendell L *et al.* 2010 Why copy others? Insights from the Social Learning Strategies Tournament. *Science* **328**, 208–213. (doi:10.1126/science.1184719)
- Laland KN. 2017 *Darwin's unfinished symphony: how culture made the human mind*. Princeton, NJ: Princeton University Press. (doi:10.23943/princeton/9780691182810.003.0001)
- Galef BG, Laland KN. 2005 Social learning in animals: empirical studies and theoretical models. *BioScience* **55**, 489. (doi:10.1641/0006-3568(2005)055[0489:sliaes]2.0.co;2)
- Morgan TJH, Rendell LE, Ehn M, Hoppitt W, Laland KN. 2012 The evolutionary basis of human social learning. *Proc. R. Soc. B* **279**, 653–662. (doi:10.1098/rspb.2011.1172)
- Hewlett BS, Fouts HN, Boyette AH, Hewlett BL. 2011 Social learning among Congo Basin hunter–gatherers. *Philos. Trans. R. Soc. B* **366**, 1168–1178. (doi:10.1098/rstb.2010.0373)
- Wang Z, Meltzoff AN, Williamson RA. 2015 Social learning promotes understanding of the physical world: Preschool children's imitation of weight sorting. *J. Exp. Child Psychol.* **136**, 82–91. (doi:10.1016/j.jecp.2015.02.010)
- Boyd R, Richerson PJ, Henrich J. 2013 The cultural evolution of technology: facts and theories. In *Cultural evolution: society, technology, language, and religion*, pp. 119–142. Cambridge, MA: The MIT Press. (doi:10.7551/mitpress/9894.003.0011)
- Basalla G. 1988 *The evolution of technology*. Cambridge, UK: Cambridge University Press.
- Richerson PJ, Boyd R. 2005 *Not by genes alone*. Chicago, IL: University of Chicago Press. (doi:10.7208/chicago/9780226712130.001.0001)
- Henrich J. 2015 *The secret of our success: how culture is driving human evolution, domesticating our species, and making us smarter*. Princeton, NJ: Princeton University Press.
- French MJ, Gravadahl J, French M. 1985 *Conceptual design for engineers*. New York, NY: Springer.
- Bonawitz E, Shafto P, Gweon H, Goodman ND, Spelke E, Schulz L. 2011 The double-edged sword of pedagogy: Instruction limits spontaneous exploration and discovery. *Cognition* **120**, 322–330. (doi:10.1016/j.cognition.2010.10.001)
- Yahosseini KS, Reijula S, Molleman L, Moussaid M. 2018 Social information can undermine individual performance in exploration-exploitation tasks. In *CogSci 2018: 40th Annual Meeting of the Cognitive Science Society: changing / minds - Monona Terrace Community and Convention Center*, Madison, WI, pp. 2473–2478. Austin, TX: Cognitive Science Society.
- Crilly N, Cardoso C. 2017 Where next for research on fixation, inspiration and creativity in design? *Des. Stud.* **50**, 1–38. (doi:10.1016/j.destud.2017.02.001)
- Derex M, Bonnefon JF, Boyd R, Mesoudi A. 2019 Causal understanding is not necessary for the improvement of culturally evolving technology. *Nat. Hum. Behav.* **3**, 446–452. (doi:10.1038/s41562-019-0567-9)
- Wiley J. 1998 Expertise as mental set: The effects of domain knowledge in creative problem solving. *Mem. Cogn.* **26**, 716–730. (doi:10.3758/bf03211392)
- Thomas C, Didierjean A, Kuhn G. 2018 It is magic! How impossible solutions prevent the discovery of obvious ones? *Q. J. Exp. Psychol.* **71**, 2481–2487. (doi:10.1177/1747021817743439)
- Ho MK, Cohen JD, Griffiths TL. 2023 Rational simplification and rigidity in human planning. *Psychol. Sci.* **34**, 1281–1292. (doi:10.1177/09567976231200547)
- Arrington CM, Logan GD. 2004 The cost of a voluntary task switch. *Psychol. Sci.* **15**, 610–615. (doi:10.1111/j.0956-7976.2004.00728.x)
- Miu E, Morgan TJH. 2020 Cultural adaptation is maximised when intelligent individuals rarely think for themselves. *Evol. Hum. Sci.* **2**, e43. (doi:10.1017/ehs.2020.42)
- Miu E *et al.* 2024 The refinement paradox and cumulative cultural evolution: collective improvement in knowledge favors conformity, blind copying and hyper-credulity. *bioRxiv*. (doi:10.1101/2024.03.22.586239)
- Morin O. 2016 *How traditions live and die*. Oxford, UK: Oxford University Press.
- Mercier H. 2020 *Not born yesterday: the science of who we trust and what we believe*. Princeton, NJ: Princeton University Press.
- Kendal RL, Boogert NJ, Rendell L, Laland KN, Webster M, Jones PL. 2018 Social learning strategies: bridge-building between fields. *Trends Cogn. Sci.* **22**, 651–665. (doi:10.1016/j.tics.2018.04.003)
- Muthukrishna M, Morgan TJH, Henrich J. 2016 The when and who of social learning and conformist transmission. *Evol. Hum. Behav.* **37**, 10–20. (doi:10.1016/j.evolhumbehav.2015.05.004)
- Rieucou G, Giraldeau LA. 2011 Exploring the costs and benefits of social information use: an appraisal of current experimental evidence. *Philos. Trans. R. Soc. B* **366**, 949–957. (doi:10.1098/rstb.2010.0325)
- Caldwell CA, Millen AE. 2010 Conservatism in laboratory microsocieties: unpredictable payoffs accentuate group-specific traditions. *Evol. Hum. Behav.* **31**, 123–130. (doi:10.1016/j.evolhumbehav.2009.08.002)
- Mathew S, Perreault C. 2015 Behavioural variation in 172 small-scale societies indicates that social learning is the main mode of human adaptation. *Proc. R. Soc. B* **282**, 20150061. (doi:10.1098/rspb.2015.0061)
- Yu H, Siegel JZ, Clithero JA, Crockett MJ. 2021 How peer influence shapes value computation in moral decision-making. *Cognition* **211**. (doi:10.1016/j.cognition.2021.104641)
- Brady WJ, McLoughlin K, Doan TN, Crockett MJ. 2021 How social learning amplifies moral outrage expression in online social networks. *Sci. Adv.* **7**, e5641. (doi:10.1126/sciadv.abe5641)
- Derex M, Boyd R. 2016 Partial connectivity increases cultural accumulation within groups. *Proc. Natl. Acad. Sci.* **113**, 2982–2987. (doi:10.1073/pnas.1518798113)
- Baer M. 2010 The strength-of-weak-ties perspective on creativity: a comprehensive examination and extension. *J. Appl. Psychol.* **95**, 592–601. (doi:10.1037/a0018761)
- Hundscheil A, Razinskas S, Backmann J, Hoegl M. 2022 The effects of diversity on creativity: a literature review and synthesis. *Appl. Psychol.* **71**. (doi:10.1111/apps.12365)
- Team RDC. 2011 *R: a language and environment for statistical computing*. Vienna, Austria: R Foundation for Statistical Computing.
- McElreath R. 2016 *Statistical rethinking: a Bayesian course with examples in R and Stan*. Boca Raton, FL: CRC Press.
- Derex M, Bonnefon JF, Boyd R, Mesoudi A. 2024 Testing the influence and robustness of socially inherited theories regarding a physical system. OSF. (doi:10.17605/OSF.IO/G74Y5)
- Derex M, Bonnefon JF, Boyd R, McElreath R, Mesoudi A. 2025 Supplementary material from: Social learning preserves both useful and useless theories by canalizing learners' exploration. Figshare. (doi:10.6084/m9.figshare.c.7632413)