

February 2025

“Social learning preserves both useful and useless theories by canalizing learners’ exploration”

Maxime Derex, Jean-François Bonnefon, Robert Boyd,
Richard McElreath and Alex Mesoudi,

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33

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, Alex Mesoudi⁶

¹ Toulouse School of Economics, CNRS, University of Toulouse Capitole, Toulouse, France.

² Institute for Advanced Study in Toulouse, University of Toulouse Capitole, Toulouse, France

³ Institute of Human Origins, Arizona State University, Tempe, AZ 85287, USA.

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

⁵ Max Planck Institute for Evolutionary Anthropology, Leipzig, Germany

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

*Correspondence to: maxime.derex@iast.fr

Abstract

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.

34 **Introduction**

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

41 Anecdotal evidence, however, suggests that learning from others can impede the
42 discovery of better alternatives. For instance, a less-than-optimal turbine blade design in the
43 early days of aircraft gas turbines went unnoticed for many years (13, p.187). Turbine blades
44 are heated by high-temperature exhaust gases resulting from the combustion of fuel in the
45 engine. An early arbitrary decision was to position the blades' fixation point near the inner end
46 of the airfoil. This caused the turbine disc, the part to which blades are attached, to overheat
47 due to heat transfer. Consequently, engineers were compelled to use a type of steel that was
48 dense, expensive, and less reliable than alternatives. It took almost a decade to resolve this
49 issue, by simply increasing the distance between the inner end of the airfoil and the blades'
50 fixation point, thereby reducing heat transfer to the turbine disc (13)).

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

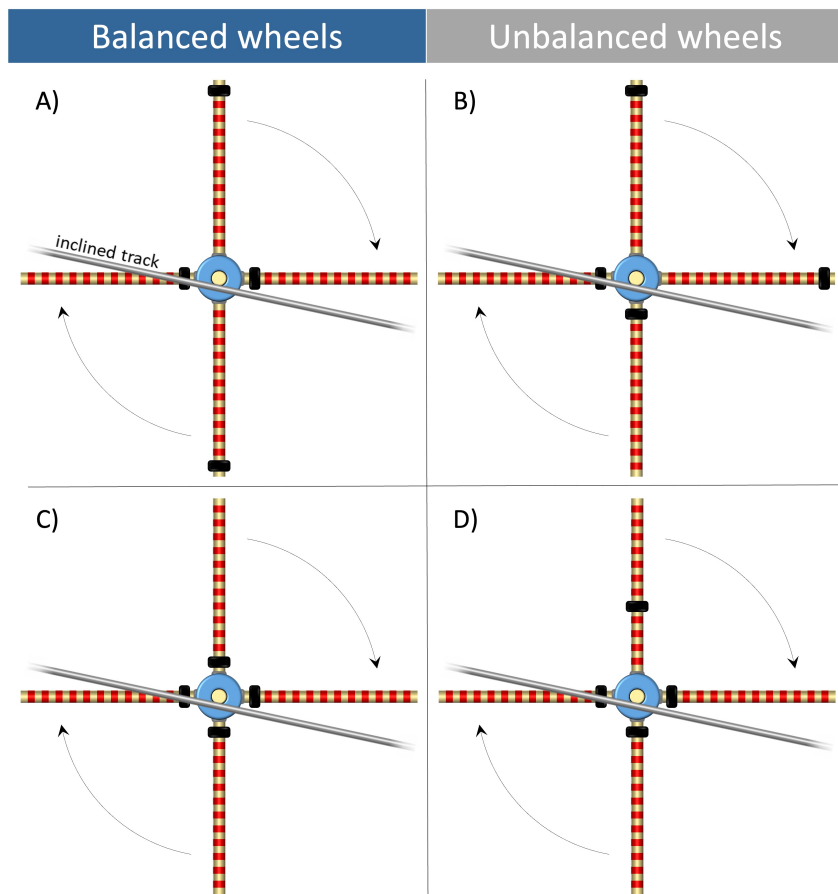
58 These results suggest that the transmission of cultural information might promote the
59 persistence of arbitrary or sub-optimal solutions by preventing social learners from thoroughly
60 exploring the solution space. However, current research is limited in several critical ways.
61 Firstly, studies often involve scenarios where learners lack objective feedback about their
62 performance and, crucially, are unable to iteratively refine their solutions. This neglects the
63 potential impact of repeated trials and objective feedback, which could help learners discover
64 more rewarding, yet initially overlooked, areas of the solution space, thereby encouraging them
65 to move beyond existing solutions (15). Secondly, existing studies are typically constrained by

66 tasks that do not allow for a detailed mapping of how social learners explore the solution space.
67 Consequently, we do not know whether social information diminishes learners' overall
68 exploration or merely channels it towards specific areas of the solution space.

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

76 Our experimental task comprises a physical system with a wheel that travels down a 1-
77 m-long inclined track, as previously used in Derex et al. (2019). The wheel has four radial
78 spokes each of which has a weight that can be moved along its length to one of 12 positions,
79 creating a space of 20,736 unique configurations (Fig. 1). The aim for the participant is to
80 position the spoke weights to minimize the time taken for the wheel to descend the track (video
81 recordings available at [this link](#)). The time it takes for the wheel to do this is determined by
82 two variables: its moment of inertia (henceforth 'Inertia') and the position of its centre of mass
83 ('CoM'). The inertia of the wheel depends on how mass is distributed around the axis. The
84 wheel has lower Inertia and will rotate more easily when weights are closer to the axis of
85 rotation. Asymmetrical wheels do not have their CoM on the axis of rotation, which can give
86 wheels better initial acceleration. These two variables imply that four qualitatively distinct
87 types of theories can be generated about the physical system: CoM theories consider CoM but
88 ignore Inertia (e.g. 'The wheel covers the track faster when its top and right weights are farther
89 from the axis than its bottom and left weights'); Inertia theories consider Inertia but ignore
90 CoM (e.g. 'The wheel covers the track faster when all its weights are close to the axis'); Correct
91 theories consider both Inertia and CoM (e.g. 'The wheel covers the track faster when all its
92 weights, except the top one, are close to the axis'); and Misleading theories consider neither
93 CoM nor Inertia (e.g. 'The wheel covers the track faster when its vertical weights are farther
94 away from the axis than the horizontal weights').

95



96

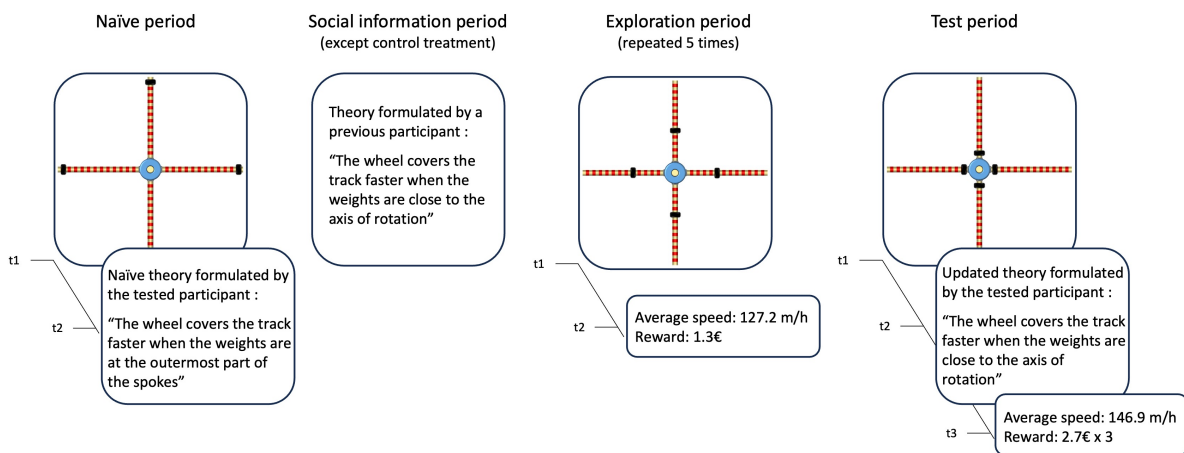
97 **Figure 1:** Illustration of the physical system used in the experiment. The wheel had four radial spokes,
 98 and one weight could be moved along each spoke. The time it takes for the wheel to cover the track is
 99 determined by its Inertia and CoM. A) A balanced wheel, consistent with the received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

110

111 Participants ($N = 200$) were exposed to a range of theories about our physical system
 112 that were generated by participants exposed to the same physical apparatus in a previous
 113 experiment (17). Participants were randomly assigned to either a Control treatment or one of
 114 four Social Information treatments (40 participants per treatment). Participants assigned to a
 115 Social Information treatment received a theory that was either Correct, Inertia-related, CoM-
 116 related, or Misleading (see Methods for details). Control participants were exposed to no

117 theory. This allowed us to study the persistence of existing theories, and study their effects on
 118 social learners' exploration patterns.

119 The experiment was preregistered ([link](#)) and organized as follows: all participants were
 120 first asked to choose a configuration, and then asked to write a theory about what makes the
 121 wheel cover the track in the shortest amount of time. During this naïve period, participants
 122 relied solely on their prior knowledge or intuition, as they had not yet observed the wheel going
 123 down the track (Fig. 2). Then, participants from the social information treatments received one
 124 of the 4 types of theory. All participants were then given the opportunity to change the
 125 configuration they had initially chosen. This was followed by an exploration period during
 126 which participants had 5 successive trials to optimize their wheel and maximize their payoff.
 127 After each trial, participants were automatically provided with their wheel's average speed and
 128 the associated payoff (range: 0-3€). After completing five trials, participants moved to the test
 129 period. They were invited to choose a bonus configuration whose associated payoff was
 130 multiplied by three (range: 0-9€) before being asked again to provide a theory about the wheel.
 131 Finally, they were provided with their final wheel's average speed and the associated payoff.



132
 133 **Figure 2:** Overview of the experimental procedure. Naïve period: Participants rely solely on their prior
 134 knowledge. They choose a configuration (t1) and then write a theory about what makes the wheel cover
 135 the track in the shortest time (t2). Social information period: All participants, except those in the control
 136 treatment, receive one of the four types of theories formulated by a participant who was exposed to the
 137 same task in a previous experiment. Exploration period: Participants interact with the physical system
 138 for five trials. At each trial, they choose a configuration (t1), the wheel is released, and participants are
 139 automatically provided with their wheel's average speed and the associated payoff (t2). Test period:
 140 Participants choose a bonus configuration, whose associated payoff is multiplied by 3 (t1), and then
 141 write a potentially updated theory about what makes the wheel cover the track in the shortest time (t2).
 142 Finally, the wheel is released, and participants are provided with their wheel's average speed and the
 143 associated payoff (t3).

144 Our main preregistered hypotheses were that: 1) social learning promotes the
145 persistence of whichever theory is received because 2) those received theories canalize
146 learners' exploration and prevent them from thoroughly exploring the solution space.

147

148 **Results**

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

155 *Understanding patterns among naïve participants*

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

166 The probabilities of producing each type of theory were comparable between
167 participants from the control treatment and participants from the social information treatments
168 prior to receiving a theory (CoM: diff. in prob. 95% CI [-0.12, 0.16], mean = 0.03; Inertia: diff.
169 in prob. 95% CI [-0.16, 0.04], mean = -0.04; Correct: diff. in prob. 95% CI [-0.02, 0.04], mean
170 = 0.01; Misleading: diff. in prob. 95% CI [-0.08, 0.01], mean = -0.02; Others : diff. in prob.
171 95% CI [-0.13, 0.18], mean = 0.02).

172 *Effect of learning in the Control treatment*

173 After interacting with our physical system, the Inertia theory was most common among
174 participants from the Control treatment (17/40 = 0.425), representing a reliable positive change
175 compared to when participants were naïve (difference in prob. 95% CI [0.15, 0.50), mean =
176 0.32). Theories categorized as Others were the second most common type produced by
177 experienced participants (16/40 = 0.40), which represented a reliable decrease compared to
178 when participants were naïve (diff. in prob. 95% CI [-0.46, -0.03), mean = -0.25). The
179 probabilities of producing the CoM, Correct, and Misleading theories were not affected by
180 interacting with the physical system (CoM: 7/40 = 0.175, diff. in prob. 95% CI [-0.23, 0.14),
181 mean = -0.05; Correct: 0/40, diff. in prob. 95% CI [-0.02, 0.02), mean = 0; and Misleading:
182 0/40, diff. in prob. 95% CI [-0.09, 0.01), mean = -0.02).

183 *Effect of social information*

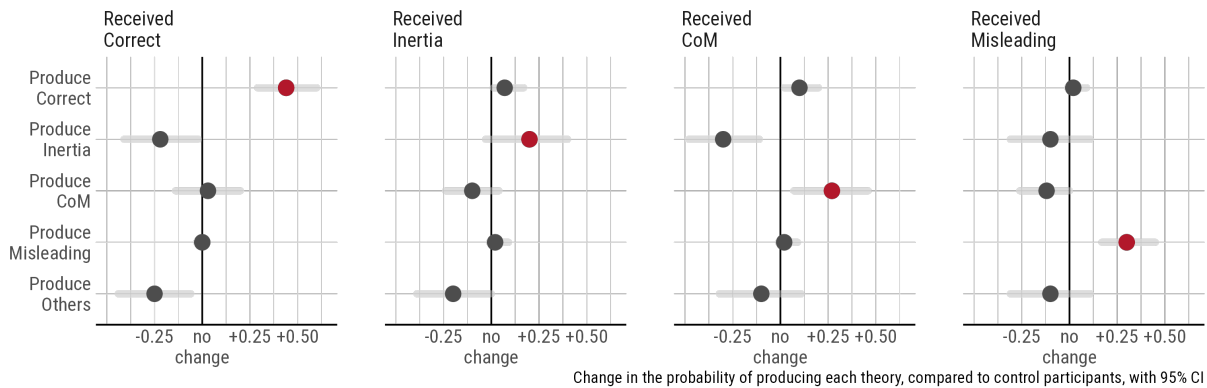
184 Participants who received a theory were more likely to change their wheel's initial
185 configuration than participants from the control treatment, indicating that participants were
186 influenced by the social information they received (Control: 0.15; CoM: 0.40, diff. in prob.
187 95% CI [0.05, 0.43], mean = 0.25; Inertia: 0.45, diff. in prob. 95% CI [0.10, 0.48], mean =
188 0.29; Correct: 0.43, diff. in prob. 95% CI [0.08, 0.45], mean = 0.27; Misleading: 0.35, diff. in
189 prob. 95% CI [0.01, 0.38], mean = 0.20). However, in accordance with our preregistered
190 hypothesis, individuals' probability of changing their initial configuration was statistically
191 comparable between social information treatments, indicating that the relevance of the received
192 theories for reaching higher performances did not affect individuals' willingness to take social
193 information into account (e.g. Misleading versus Correct: diff. in prob. 95% CI [-0.13, 0.28],
194 mean = 0.07). Contrary to one of our secondary preregistered hypotheses, individuals'
195 probability of changing their initial configuration was not lower in men compared to women
196 (Men: 0.43; Women: 0.39; diff. in prob. 95% CI [-0.18, 0.12], mean = -0.04).

197 We now look at the effect of interacting with the task among participants who received
198 social information. Our results confirm our preregistered hypothesis that receiving a theory
199 about how a task works before interacting with the task increases the likelihood that individuals
200 will produce the same theory after interacting with the task, compared to participants who do
201 not receive a theory (Fig. 3). The probability of producing the Correct theory after interacting
202 with the task was 0 among participants from the Control treatment and 0.45 among participants
203 who received the correct theory, which represents a reliable increase compared to participants
204 from the Control treatment (diff. in prob. 95% CI [0.29, 0.59], mean = 0.44). Reliable increases

205 in the probability of producing the theory received were also observed when participants
 206 received the Inertia theory (from 0.43 to 0.63, diff. in prob. 95% CI [-0.01, 0.41], mean = 0.20),
 207 the CoM theory (from 0.18 to 0.45, diff. in prob. 95% CI [0.07, 0.45], mean = 0.27) and the
 208 Misleading theory (from 0 to 0.30, diff. in prob. 95% CI [0.17, 0.44], mean = 0.30).

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.



209

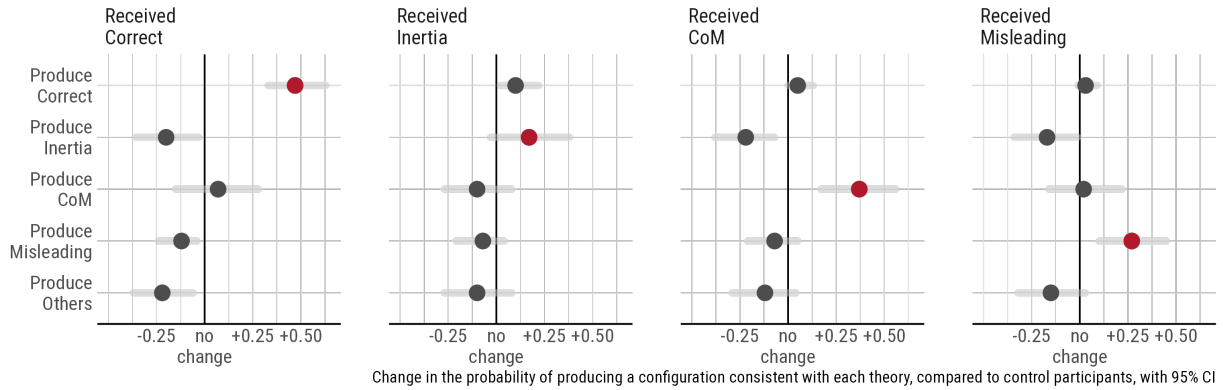
210 **Figure 3:** Difference in the probability of producing each type of theory as a function of the theory
 211 received, compared to receiving no theory. Participants in the Control treatment (no theory) produced
 212 one out of five types of theories after interacting with our physical system: Correct (0%), Inertia
 213 (42.5%), CoM (17.5%), Misleading (0%), Others (40%). The figure illustrates the difference in the
 214 probability of producing each type of final theory as a function of the theory received, compared to
 215 those values. For instance, the left column illustrates the difference in the probability of producing each
 216 type of theory between participants in the Control treatment and participants who received the Correct
 217 theory. When participants received the Correct theory, the probability of producing each type of theory
 218 changed as follows: Correct went from 0 to 0.45 (+0.45), Inertia from 0.42 to 0.20 (-0.22), CoM from
 219 0.17 to 0.20 (+0.03), and Misleading from 0 to 0 (0), respectively. Red dots along the diagonal indicate
 220 that, compared to participants in the Control treatment, receiving any theory before interacting with the
 221 task increases individuals' probability of producing the same theory after completing the task.

222

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

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.



230

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

236

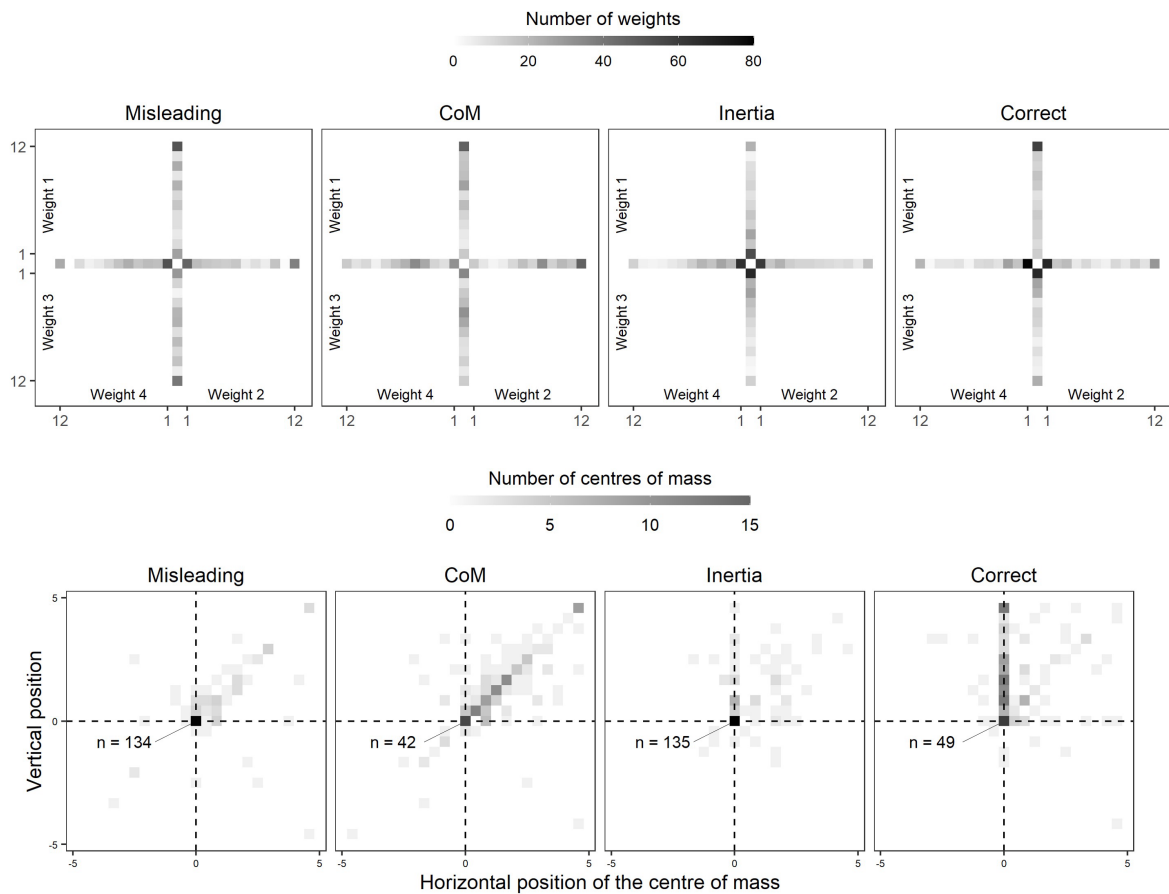
237 Among participants in the Correct treatment, the probability of writing a theory
238 categorized as Correct was comparable to the probability of producing a configuration
239 consistent with the Correct theory (0.45 and 0.47, respectively, diff. in prob. 95% CI [-0.23,
240 0.18], mean = -0.03). This contradicts our secondary preregistered hypothesis that correct, two-
241 dimensional theories might fail to be properly transmitted through text messages because
242 participants would tend to write simpler, unidimensional theories that are less costly to
243 articulate.

244 *Effect of social information on exploration patterns*

245 Over the 5 trials of the exploration period, participants in the Control treatment produced an
246 estimated mean of 4.75 unique configurations (95% CI [4.57, 4.94]), which is comparable to
247 the levels of exploration exhibited by participants who received a theory (contrasts with:
248 Misleading: 95% CI [-0.37, 0.12], mean = -0.13; CoM: 95% CI [-0.20, 0.30], mean = 0.05;
249 Inertia: 95% CI [-0.35, 0.15], mean = -0.10; Correct: 95% CI [-0.20, 0.30], mean = 0.05).

250 However, consistent with our preregistered hypothesis, exploration patterns were
251 influenced by the type of theory received (Fig. 5). Participants who received the Misleading
252 theory mostly produced balanced wheels with their horizontal weights closer to the axis than
253 vertical weights (Fig. 5.a,e). Participants who received the CoM theory mostly produced
254 unbalanced wheels with their top and right weights at the outermost position (Fig. 5.b,f).

255 Participants who received the Inertia theory mostly produced balanced wheels with all their 4
 256 weights close to the axis (Fig. 5.c,g). Participants who received the Correct theory mostly
 257 produced unbalanced wheels with all their weights, except the top one, close to the axis (Fig.
 258 5.d,h).



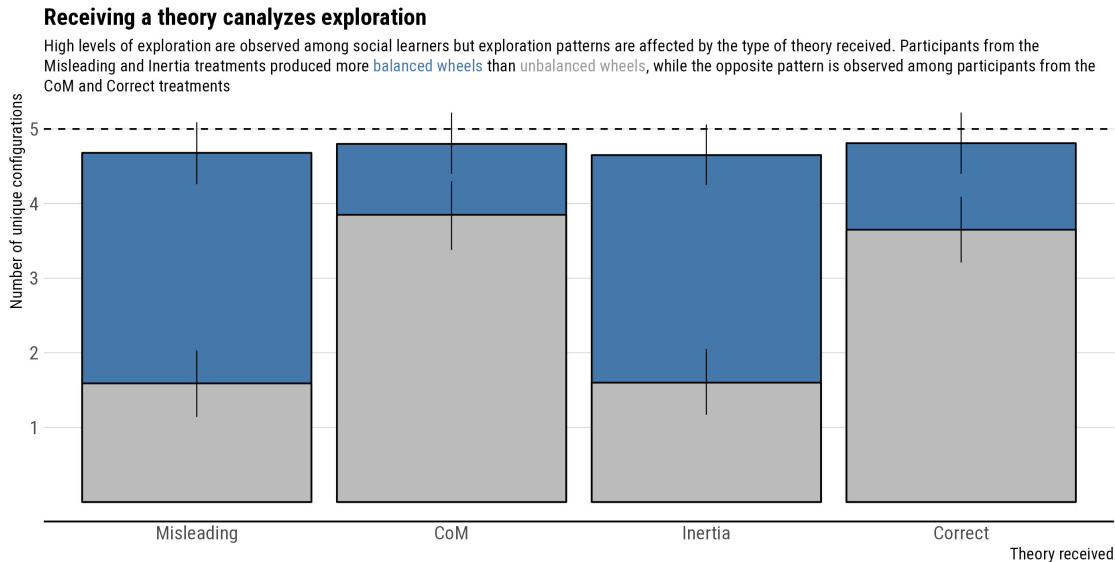
259

260 **Figure 5:** Exploration patterns are affected by the type of theory received. The top panel shows heat
 261 maps illustrating the most frequent weight positions along each spoke in each treatment. The bottom
 262 panel shows the most frequent positions of the wheels' centre of mass in each treatment. Wheels with
 263 both their horizontal and vertical weights equidistant from the axis are balanced, with their centre of
 264 mass located on the axis of rotation (i.e. at the intersection of the dotted lines in the bottom panel).
 265 Values 1–12 in the top panel describe the positions of weights 1–4. Values –5 to +5 in the bottom panel
 266 describe the x- and y-coordinates of the wheels' centre of mass.

267

268 Exploratory analyses indicate that participants in the Misleading and Inertia treatments
 269 reliably produced more balanced wheels (Misleading: 95% CI [2.68, 3.50], mean = 3.09;
 270 Inertia: 95% CI [2.64, 3.45], mean = 3.05) compared to unbalanced wheels (Misleading: 95%
 271 CI [1.18, 2.02], mean = 1.60; Inertia: 95% CI [1.16, 2.04], mean = 1.60). This suggests a greater
 272 focus on varying the moment of inertia of the wheel rather than the position of its centre of
 273 mass. In the CoM and Correct treatments, the opposite exploration pattern was observed.

274 Participants in these treatments reliably produced fewer balanced wheels (CoM: 95% CI [0.55,
 275 1.37], mean = 0.95; Correct: 95% CI [0.74, 1.56], mean = 1.16) compared to unbalanced wheels
 276 (CoM: 95% CI [3.41, 4.29], mean = 3.85; Correct: 95% CI [3.21, 4.08], mean = 3.65),
 277 indicating a greater focus on varying the position of the wheel's centre of mass rather than its
 278 moment of inertia.



279

280 **Figure 6:** Receiving a theory affects the exploration of dimensions relevant to the performance of the
 281 wheel. Errors bars show 95% Confidence Intervals. The horizontal line indicates the maximum number
 282 of unique configurations that can be produced during the exploration period. (See Figure 1 for examples
 283 of balanced and unbalanced wheels).

284

285 *Effect of social information on performance*

286 Participants in the Inertia and Correct treatments performed better in the bonus trial than
 287 those in the Control treatment (Inertia 95% CI [0.00, 1.86], mean = 0.93; Correct 95% CI [0.00,
 288 1.85], mean = 0.92). In contrast, participants in the CoM and Misleading treatments performed
 289 comparably to the Control group (CoM 95% CI [-1.55, 0.28], mean = -0.63; Misleading 95%
 290 CI [-1.16, 0.70], mean = -0.22). This partially contradicts our preregistered hypothesis that
 291 partially- and fully-correct theories would be beneficial, while the misleading theory would be
 292 detrimental.

293 Expanding the analysis to all trials reveals that participants improved their performance
 294 across trials in all treatments (Misleading : 95% CI [0.09, 0.17], mean = 0.13; CoM: 95% CI
 295 [0.15, 0.20], mean = 0.17; Inertia: 95% CI [0.15, 0.21], mean = 0.18; Correct: 95% CI [0.15,
 296 0.22], mean = 0.18). Across all trials, the performance of participants who received a theory

297 differed from that of participants in the Control treatment, with the exception of the CoM
298 treatment (contrast with Control treatment 95% CI [-0.17, 0.16], mean = -0.01; Trial ×
299 Treatment 95% CI [-0.05, 0.02], mean = -0.01). Both the Inertia and Correct treatments resulted
300 in better performance (Inertia: contrast with Control treatment 95% CI [0.05, 0.27], mean =
301 0.16; Trial × Treatment 95% CI [-0.03, 0.01], mean = -0.01; Correct: contrast with Control
302 treatment 95% CI [0.01, 0.18], mean = 0.10; Trial × Treatment 95% CI [-0.02, 0.01], mean =
303 0.00). Participants who received the Misleading theory displayed slower improvement than
304 participants from the Control treatment (contrast with Control treatment 95% CI [-0.14, 0.49],
305 mean = 0.18; Trial × Treatment 95% CI [-0.12, 0.01], mean = -0.05). Additional analyses
306 indicate that the relatively poor performance of participants in the CoM treatment results from
307 a reliably higher probability of producing wheels that did not complete a full first revolution
308 due to extreme positions of the wheel's centre of mass (Supplementary Table 1).

309

310 **Discussion**

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

320 These results support recent experimental findings highlighting what has been termed
321 the double-edged sword of pedagogy, where teaching increases the likelihood of performing
322 relevant behaviours but also reduces the likelihood of discovering alternative ones (14).
323 However, our study offers a different perspective on how social information affects learners'
324 exploration. While previous research has suggested that social learning limits exploration, our
325 results indicate that social learning canalizes learners' exploration without necessarily reducing
326 its overall extent. The tasks used in earlier experiments may have masked similar effects. For
327 instance, in Bonawitz et al.'s study, children were given unfamiliar toys with four non-obvious
328 functions, such as squeaking, flashing lights and a hidden mirror. Children who were informed

329 about one function of the toy were less likely to discover the other functions compared to those
330 who were not given any information(14). However, the study also found that children who
331 were told about a function spent more time using that demonstrated function. This increased
332 time spent was interpreted as evidence of limited exploration. This interpretation assumes that
333 discovering only the different functions was the sole objective. Yet, it seems reasonable to
334 assume that by spending more time using the demonstrated function, children might have
335 discovered additional properties of that function (for instance, how to modulate the intensity
336 or tone of the squeaking function).

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

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

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

377 Our results contribute to the debate on humans' propensity to rely on social information
378 and the implications for the persistence of arbitrary solutions. Some scholars argue that
379 individuals are inclined to rely on social information, which can facilitate the adoption of hard-
380 to-devised, unintuitive solutions (12, 22, 23). Others contend that individuals should be cautious
381 about social information to avoid the risk of being accidentally or intentionally misled (24, 25).
382 According to the former view, arbitrary cultural solutions are likely to persist, while the latter
383 suggests these solutions should quickly fade unless they are particularly intuitive to individuals.
384 Our experiment supports the former view, challenging the idea that arbitrary solutions persist
385 only when they are intuitive to participants. Notably, only one out of 40 participants produced
386 a theory categorized as Misleading during the naïve period, and none produced a theory
387 categorized as Correct. Despite this, both types of theory persisted in our experiment.

388 One might argue that our experimental design was biased towards observing a persistent
389 impact of social information. Indeed, theory predicts that individuals are likely to heavily rely
390 on social information when they are uncertain (26), either because they have no relevant prior
391 information (1), because the number of potential solutions is large (27), or because others
392 possess more reliable information (28). In our experiment, participants faced an unfamiliar task

393 with a large solution space and were informed that the theories they received came from
394 individuals with prior experience with the physical system. However, we believe these
395 conditions appropriately reflect the challenges individuals face when attempting to improve
396 existing technologies. Moreover, at least two aspects of our design make the experiment
397 conservative. First, our task was low-dimensional (i.e. only Inertia and CoM affect the
398 dynamics of the wheel), which is likely to reduce uncertainty, and thereby reliance on social
399 information, compared to facing a real technology, which tends to be high-dimensional.
400 Second, we used non-noisy payoffs, whereas most real technologies provide noisy payoffs,
401 known to increase the persistence of arbitrary solutions (29).

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

424

425 **Methods**

426 Participants

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

432 Ethical statement

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

438 Experimental apparatus

439 The experimental apparatus was similar to that used in (17). It consisted of a wheel that had to
440 travel down a 1-m-long inclined track. The wheel had four radial spokes, and one weight could
441 be moved along each spoke. Weights could be placed on one of 12 discrete positions which
442 created a space of 20,736 unique configurations. The performance of the wheel depends on two
443 variables: its moment of inertia and the position of its centre of mass. The wheel's moment of
444 inertia depends on how mass is distributed around its axis of rotation. Wheels with a smaller
445 moment of inertia (i.e. wheels that have their weights closer to the axis) require less torque to
446 increase angular momentum and spin faster (Video recordings are available at [this link](#)).
447 Asymmetrical wheels do not have their centre of mass on the axis of rotation, which can
448 provide a better initial acceleration. When the centre of mass of the wheel is in the wheel's
449 upper right quadrant (assuming the wheel goes downhill from left to right), more potential
450 energy is converted into angular kinetic energy so that the wheel will benefit from higher
451 increases in angular momentum. In our experiment, both the wheel's moment of inertia and the
452 position of its centre of mass had to be taken into account to reach the best performance. A
453 higher centre of mass can produce better acceleration, but it will increase the wheel's moment
454 of inertia and so there was a tradeoff between maximizing acceleration and minimizing inertia
455 (17).

456 Procedure

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

465 Experimental design

466 The experiment comprised 3 distinct periods: a naïve period, an exploration period and a test
467 period. Participants chose their configurations through a computer program using 4 sliders.

468 *Naïve period (Trial 0)*

469 Participants started by choosing a configuration. Right after participants confirmed their
470 configuration, they were asked to write their theory about what makes the wheel covering the
471 distance in the shortest amount of time. Theories had to be less than 340 characters long and
472 always started with ‘The wheel covers the distance faster when...’ in order to encourage
473 participants to provide a general statement about the wheel. During this period, participants
474 relied solely on their prior knowledge as they had not yet observed the wheel being released.
475 Participants then received one of the 5 experimental treatments before being given the
476 opportunity to change the configuration they had initially chosen.

477 *Exploration period (Trials 1-5)*

478 Once participants confirmed the configuration of their wheel (whether they had changed it or
479 not), the experimenter positioned the weights on the physical wheel accordingly (the computer
480 screen was projected onto a wall to the right of the participant in order to allow the experimenter
481 to see the chosen configuration without interacting with the participant). The wheel was then
482 positioned on the rails. A mechanical lever maintained the wheel motionless, with 2 of its
483 spokes parallel to the ground at its starting position. Once released, the time it took the wheel
484 to descend the track was automatically recorded by the computer program, and the wheel’s
485 average speed and associated payoff was automatically displayed on the participant’s screen.

486 The participant could then choose a new configuration. The procedure was repeated until
487 participants had completed their 5 trials.

488 *Test period (Bonus configuration)*

489 After having completed their 5 trials, participants were invited to choose a bonus configuration
490 whose associated score is multiplied by 3. After confirming their configuration, participants
491 were asked to provide a theory about the wheel. Only after providing their theory, participants
492 could observe the wheel going down.

493 *Experimental treatments*

494 Five treatments were run. All participants except those in the control treatment were provided
495 with social information. Social information took the form of a theory that was produced by a
496 participant from a previous experiment involving the same physical apparatus. Four theories
497 that vary in their accuracy were chosen in order to cover the three qualitatively distinct types
498 of theories that participants can receive: fully correct, partially correct and misleading. In the
499 Correct treatment, participants were provided with a theory that states: *'The wheel covers the*
500 *distance faster when all weights are close to the axis except the top weight that has to be a bit*
501 *farther away'*. This theory encourages participants to produce wheels with a low moment of
502 inertia and a centre of mass located above the wheel's axis of rotation (at the wheel's initial
503 position). This theory captures the two principles that allow participants to produce the most
504 efficient wheels. In the Inertia treatment, participants were provided with a partially correct
505 theory that emphasizes solely the role of the wheel's moment of inertia: *'The wheel covers the*
506 *distance faster when the weights are close to the axis'*. In the Centre of Mass treatment,
507 participants were provided with another partially correct theory, the one that emphasizes solely
508 the role of the position of the wheel's CoM: *'The wheel covers the distance faster when the top*
509 *and right weights are farther from the axis than the bottom and left weights'*. In the Misleading
510 treatment, participants were provided with a theory that does not emphasize either of the two
511 dimensions that are relevant to the performance of the wheel: *'The wheel covers the distance*
512 *faster when the horizontal weights are closer to the axis than the vertical weights'*. This
513 misleading theory was chosen because 1) it produces a recognizable wheel configuration that
514 does not overlap with configurations consistent with the other seeded theories, ensuring we
515 could identify the effect of social information and 2) it was rare in the pre-existing dataset,
516 minimizing the likelihood that its "stickiness" was due to participants' prior intuitions rather
517 than the influence of social information.

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

523 *Pre-experiment information*

524 Instructions could be read on a computer screen and stated that the participants' task was to
525 position 4 weights on a wheel in order to minimize the time it takes the wheel to cover an
526 inclined track. Participants were informed that they will have 5 trials to do this and that their
527 payoff will be determined by the performance of each of their wheels. Additionally, they were
528 told that they will have to formulate a theory about the wheel and that this theory will be
529 evaluated and will determine part of their score. Participants were informed that they might be
530 provided with a previous participant's theory and that this theory might help them maximizing
531 their score.

532 *Participants' payoff*

533 During the exploration period, the following equation determined the payoff of each wheel:

$$534 \quad [1 - ((\text{MaxSpeed} - \text{RecordedSpeed}) / (\text{MaxSpeed} - \text{MinSpeed}))] \times 3 + \text{Bonus}$$

535 with $\text{MaxSpeed} = 160$, $\text{MinSpeed} = 96$. RecordedSpeed was the recorded average speed of the
536 wheel. Bonus took the value 0.2 for wheels that descended the rails and 0 otherwise. During
537 the test period, wheels' payoff was multiplied by 3. Theories formulated by participants were
538 immediately evaluated by the experimenter (but later independently coded for the purposes of
539 statistical analyses, see below) and provided participants with an additional payoff of 0, 0.5 or
540 1 euro depending on whether they mention none, one or two of the dimensions that are relevant
541 to the performance of the wheel. Participants' final payoff corresponded to the sum of the
542 payoff of each of their wheels plus the payoffs associated with their theories.

543 *Theory coding (human raters)*

544 For the purposes of statistical analyses, theories were coded by 2 sets of 3 individuals blind to
545 the research question. Coders were explained the dynamics of the wheel (i.e. the respective role
546 of the inertia and centre of mass in the performance of the wheel) before completing their task.
547 The first set of coders were asked to code participants' theories according to whether they

548 contain accurate information related to the moment of inertia (Inertia) and/or centre of mass
549 (CoM). A theory contained information related to the moment of inertia when it says that the
550 wheels goes faster when its weights are close to the axis (e.g. '*The wheel covers the track faster*
551 *when its weights are balanced and close to the axis.*'). A theory contained information related
552 to the centre of mass when it says that the wheel goes faster when its centre of mass is in the
553 upper-right quadrant (e.g. '*The wheel covers the track faster when its top and right weights are*
554 *far from the axis and its bottom and left weights are close to it.*'). Theories that contained
555 accurate information about the effect of the Inertia of the wheel were considered as partially
556 correct and consistent with the seeded Inertia theory. Theories that contained accurate
557 information about the effect of the CoM of the wheel were considered as partially correct and
558 consistent with the seeded CoM theory. Theories that contained accurate information about
559 both effects were considered as correct and consistent with the seeded Correct theory. The
560 second set of coders evaluated whether theories were consistent with the seeded Misleading
561 theory. Theories that were not categorized as Correct, CoM, Inertia or Misleading were
562 categorized as Others. A majority among coders determined final coding. Cohen's kappa
563 coefficients reveal either substantial or almost perfect agreement between raters (0.84 for
564 Inertia, 0.78 for CoM and 0.96 for Misleading).

565 *Theory coding (classification algorithm)*

566 Participants' configuration that was chosen just before formulating their updated theory (i.e.
567 the bonus trial associated with the test period) was used to infer participants' theories at the
568 end of the experiment. The most compact configuration was considered as evidence for
569 partially correct and 'Inertia-related' theory. Configurations with their centre of mass in the top
570 right quadrant, minus those that are better than the most compact wheel, were considered as
571 evidence for partially correct and 'CoM-related' theory. All configurations that had their
572 bottom and left weight at the closest position to the axis and that were better than the most
573 compact wheel, were considered as evidence for accurate theory (Supplementary Table 2). All
574 configurations that had their horizontal weights closer to the axis than their vertical weights
575 (minus those that correspond to either the CoM, Inertia or Accurate theories) were considered
576 as evidence for Misleading theory. Other configurations were categorized as Others.

577 Statistical analyses and models output

578 We ran a series of Bayesian models in R (36). Models were fitted using the *rethinking* package
579 (37) and 95% credible intervals were used to make inferences. Full details about our models
580 are available at [this link](#).

581 *Probability of changing configuration after being exposed to social information.*

582 Configurations from the naïve period and the first trial of the exploration period were used. We
583 fitted a logistic regression with ‘Change’ as the response variable and one dummy variable for
584 each social information treatment as predictor variables.

585 *Probability of changing configuration among men and women.*

586 Configurations from the naïve period and the first trial of the exploration period were used. We
587 fitted a logistic regression with ‘Change’ as the response variable and one dummy variable for
588 ‘Women’ as predictor variable.

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

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

599 *Effects of social information on the theory produced.*

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

602 *Effect of social information on the number of unique configurations produced.*

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

606 *Effect of social information on the number of balanced and unbalanced configurations.*

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

611 *Effect of social information on performance*

612 Configurations from the exploration and test periods were used. We fitted a linear model with
613 ‘Payoff’ as the outcome variable, ‘Trial’, ‘Treatment’ and ‘Trial:Treatment’ as predictor
614 variables, and ‘Participant’s identity’ as random effect.

615 *Deviation from pre-registered analyses*

616 The pre-registered analysis on the effect of social information on performance uses ‘Payoff’ as
617 outcome variable instead of ‘Speed’. ‘Payoff’ and ‘Speed’ are linearly correlated, but using
618 ‘Payoff’ as the outcome variable allows the model to sample more efficiently and provides
619 narrower confidence intervals (see Supplementary Material in (17)). This is because the gap
620 between the wheels that did not descend the rails and those that did is proportionally smaller
621 when considering Payoff (0 to 0.3€–3€) than Speed (0 to 98.2–154.7 m/h).

622 Pre-registration

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

624 **Data availability**

625 The data that support the findings of this study are available at [this link](#).

626 **Code availability**

627 Codes used in this paper are available at [this link](#).

628 **Acknowledgments**

629 M.D. acknowledges Institute for Advanced Study in Toulouse (IAST) funding from the
630 French National Research Agency (ANR) under grant no. ANR-17-EURE-0010
631 (Investissements d’Avenir programme) and ANR funding through the ANR-21-CE28-0019-
632 01 OPTILEARN project. We thank the editor and two anonymous reviewers for valuable
633 comments on a previous version.

634

635

636 **References**

- 637 1. Boyd R, Richerson PJ. Culture and the evolutionary process. Chicago : University of
638 Chicago Press; 1985.
- 639 2. Boyd R, Richerson PJ, Henrich J. The cultural niche: Why social learning is essential
640 for human adaptation. *Proc Natl Acad Sci U S A*. 2011;108:10918-25.
- 641 3. Rendell L, Boyd R, Cownden D, Enquist M, Eriksson K, Feldman MW, et al. Why
642 Copy Others? Insights from the Social Learning Strategies Tournament. *Science*.
643 2010;328(5975):208-13.
- 644 4. Laland KN. Darwin's Unfinished Symphony: How Culture Made the Human Mind:
645 Princeton University Press; 2017.
- 646 5. Galef BG, Laland KN. Social learning in animals: Empirical studies and theoretical
647 models. *Bioscience*. 2005;55(6):489-99.
- 648 6. Morgan TJH, Rendell LE, Ehn M, Hoppitt W, Laland KN. The evolutionary basis of
649 human social learning. *Proc R Soc B-Biol Sci*. 2012;279(1729):653-62.
- 650 7. Hewlett BS, Fouts HN, Boyette AH, Hewlett BL. Social learning among Congo Basin
651 hunter-gatherers. *Philosophical Transactions of the Royal Society B: Biological Sciences*.
652 2011;366(1567):1168-78.
- 653 8. Wang Z, Meltzoff AN, Williamson RA. Social learning promotes understanding of
654 the physical world: Preschool children's imitation of weight sorting. *Journal of Experimental*
655 *Child Psychology*. 2015;136:82-91.
- 656 9. Boyd R, Richerson PJ, Henrich J. The cultural evolution of technology: facts and the-
657 ories. *Cultural evolution: society, technology, language, and religion*. 2013:119-42.
- 658 10. Basalla G. The evolution of technology. Cambridge: UK: Cambridge University
659 Press; 1988.
- 660 11. Richerson PJ, Boyd R. Not by genes alone. Chicago, IL: University of Chicago Press;
661 2005.
- 662 12. Henrich J. The secret of our success: how culture is driving human evolution, domes-
663 ticating our species, and making us smarter: Princeton University Press; 2015.
- 664 13. French MJ, Gravdahl J, French M. Conceptual design for engineers: Springer; 1985.
- 665 14. Bonawitz E, Shafto P, Gweon H, Goodman ND, Spelke E, Schulz L. The double-
666 edged sword of pedagogy: Instruction limits spontaneous exploration and discovery. *Cogni-
667 tion*. 2011;120(3):322-30.
- 668 15. Yahosseini KS, Reijula S, Molleman L, Moussaid M. Social information can under-
669 mine individual performance in exploration-exploitation tasks. 2018.
- 670 16. Crilly N, Cardoso C. Where next for research on fixation, inspiration and creativity in
671 design? *Design Studies*. 2017;50:1-38.
- 672 17. Derex M, Bonnefon J-F, Boyd R, Mesoudi A. Causal understanding is not necessary
673 for the improvement of culturally evolving technology. *Nature Human Behaviour*.
674 2019;3(5):446-52.
- 675 18. Wiley J. Expertise as mental set: The effects of domain knowledge in creative prob-
676 lem solving. *Memory & Cognition*. 1998;26(4):716-30.
- 677 19. Thomas C, Didierjean A, Kuhn G. It is magic! How impossible solutions prevent the
678 discovery of obvious ones? *Quarterly Journal of Experimental Psychology*.
679 2018;71(12):2481-7.
- 680 20. Ho MK, Cohen JD, Griffiths T. Rational simplification and rigidity in human plan-
681 ning. *Psychol Sci*. 2023;34(11):1281 -92.
- 682 21. Arrington CM, Logan GD. The Cost of a Voluntary Task Switch. *Psychol Sci*.
683 2004;15(9):610-5.

- 684 22. Miu E, Morgan TJH. Cultural adaptation is maximised when intelligent individuals
685 rarely think for themselves. *Evolutionary Human Sciences*. 2020;2:e43.
- 686 23. Miu E, Rendell L, Bowles S, Boyd R, Cownden D, Enquist M, et al. The refinement
687 paradox and cumulative cultural evolution: collective improvement in knowledge favors con-
688 formity, blind copying and hyper-credulity. *bioRxiv*. 2024:2024.03.22.586239.
- 689 24. Morin O. *How traditions live and die*: Oxford University Press; 2016.
- 690 25. Mercier H. *Not born yesterday: The science of who we trust and what we believe*:
691 Princeton University Press; 2020.
- 692 26. Kendal RL, Boogert NJ, Rendell L, Laland KN, Webster M, Jones PL. Social Learn-
693 ing Strategies: Bridge-Building between Fields. *Trends Cogn Sci*. 2018;22(7):651-65.
- 694 27. Muthukrishna M, Morgan TJH, Henrich J. The when and who of social learning and
695 conformist transmission. *Evol Hum Behav*. 2016;37(1):10-20.
- 696 28. Rieucou G, Giraldeau LA. Exploring the costs and benefits of social information use:
697 an appraisal of current experimental evidence. *Philos Trans R Soc B-Biol Sci*.
698 2011;366(1567):949-57.
- 699 29. Caldwell CA, Millen AE. Conservatism in laboratory microsocieties: unpredictable
700 payoffs accentuate group-specific traditions. *Evol Hum Behav*. 2010;31(2):123-30.
- 701 30. Mathew S, Perreault C. Behavioural variation in 172 small-scale societies indicates
702 that social learning is the main mode of human adaptation. *Proceedings of the Royal Society*
703 *B: Biological Sciences*. 2015;282(1810).
- 704 31. Yu H, Siegel JZ, Clithero JA, Crockett MJ. How peer influence shapes value compu-
705 tation in moral decision-making. *Cognition*. 2021;211:104641.
- 706 32. Brady WJ, McLoughlin K, Doan TN, Crockett MJ. How social learning amplifies
707 moral outrage expression in online social networks. *Science Advances*. 2021;7(33):eabe5641.
- 708 33. Derex M, Boyd R. Partial connectivity increases cultural accumulation within groups.
709 *Proceedings of the National Academy of Sciences*. 2016;113(11):2982-7.
- 710 34. Baer M. The strength-of-weak-ties perspective on creativity: a comprehensive exami-
711 nation and extension. *Journal of applied psychology*. 2010;95(3):592.
- 712 35. Hundschell A, Razinskas S, Backmann J, Hoegl M. The effects of diversity on crea-
713 tivity: A literature review and synthesis. *Applied Psychology*. 2022;71(4):1598-634.
- 714 36. Team RDC. *R: A Language and Environment for Statistical Computing*. Vienna, Aus-
715 tria: R Foundation for Statistical Computing; 2011.
- 716 37. McElreath R. *Statistical Rethinking: A Bayesian Course with Examples in R and*
717 *Stan*: CRC Press; 2016.