

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16

Identifying innovation in laboratory studies of cultural evolution:
rates of retention and measures of adaptation

Christine A. Caldwell^{1,*}, Hannah Cornish¹, and Anne Kandler²

¹Psychology, School of Natural Sciences, University of Stirling,
Stirling, FK9 4LA, United Kingdom

²Department of Mathematics, City University London, Northampton Square,
London, EC1V 0HB, United Kingdom

*Author for correspondence c.a.caldwell@stir.ac.uk

17 Summary

18

19 In recent years, laboratory studies of cultural evolution have become increasingly prevalent
20 as a means of identifying and understanding the effects of cultural transmission on the form
21 and functionality of transmitted material. The data sets generated by these studies may
22 provide insights into the conditions encouraging, or inhibiting, high rates of innovation, as
23 well as the effect that this has on measures of adaptive cultural change. Here we review
24 recent experimental studies of cultural evolution with a view to elucidating the role of
25 innovation in generating observed trends. We first consider how tasks are presented to
26 participants, and how the corresponding conceptualisation of task success is likely to
27 influence the degree of intent underlying any deviations from perfect reproduction. We then
28 consider the measures of interest used by the researchers to track the changes that occur as a
29 result of transmission, and how these are likely to be affected by differing rates of retention.
30 We conclude that considering studies of cultural evolution from the perspective of innovation
31 provides valuable insights which help to clarify important differences in research designs,
32 which have implications for the likely effects of variation in retention rates on measures of
33 cultural adaptation.

34

35 Keywords: cultural evolution; iterated learning; micro-society; social learning; transmission
36 chain

37

38 1. Introduction

39

40 In the current article, we consider what we can learn about innovation from experimental
41 studies of cultural evolution. Here we define as *cultural* any traits (behavioural,
42 psychological, or artefactual) that exhibit heritability as a result of learning from others, with
43 *cultural evolution* referring to a process entailing modification to cultural traits over time. We
44 also refer to *cultural change* to indicate the aggregate effect of the process of cultural
45 evolution on cultural traits between particular time points.

46

47 Following these definitions, it is clear that understanding innovation is fundamental to
48 understanding cultural change. It is widely acknowledged that both innovation and social
49 learning are the two cornerstones of cultural evolution [1]. While faithful social learning (i.e.
50 social learning without any source of error) operates to maintain cultural traditions, on its
51 own it will produce only cultural stasis. It is innovation which drives cultural change.

52 Understanding the contexts which promote innovation, and the effect this has on population-
53 level shifts in behaviour, is therefore essential to understanding phenomena as diverse as
54 developments in science and technology, the rise and fall of fads and fashions, and shifting
55 societal trends.

56

57 There are now numerous experiments reported in the literature which purport to capture
58 aspects of cultural evolution under laboratory conditions. Potentially, these should offer
59 fertile ground for helping us understand the catalysts and consequences of innovation within
60 populations of learners. To our knowledge, none of these studies have been designed with the
61 explicit intention of investigating innovation, as they are more concerned with documenting
62 overall patterns of change, rather than identifying particular individuals, or particular

63 individual decisions, as the source of such change. However, some studies do provide an
64 insight into factors affecting rates of innovation. Furthermore, in studies which permit
65 inferences about variation in innovation rate, it is also possible to consider the effect this has
66 on the measures of directional cultural change used by the researcher. Although cultural
67 change requires innovations, it does not necessarily follow that high innovation rates generate
68 pronounced cultural change, aggregated over multiple learners. Depending on the
69 circumstances under consideration, innovations may not necessarily modify cultural traits in
70 consistent directions, generating limited change at the group level. In this article we focus on
71 experimental research on cultural evolution with the aim to review what we can infer about
72 the role of innovation in these studies.

73

74 1.1 Experimental studies of cultural evolution

75

76 Although experimental studies of cultural evolution may take a variety of forms, we believe
77 that all designs share certain unifying features which are worth outlining here. Firstly, in
78 contrast to more typical psychological experiments which concern how a single individual
79 performs on a task, or sometimes how one individual learns from another, in studies of
80 cultural evolution a single replicate within an experiment consists of multiple (three or more)
81 participants. In this way these designs capture the *repeated* occurrences of social learning
82 involved in cultural change, as opposed to one-off cases of individual learning or social
83 learning in general. Secondly, within each replicate, participants have some form of access to
84 information about the solutions or responses of other members of the same replicate. The
85 exact nature of the information available may vary, but can include direct observation, verbal
86 report, or stored information about solutions or responses presented remotely, i.e. in the
87 absence of their progenitor. Finally, all studies involve a measure that is repeated

88 successively, the overall aim being to describe the nature and/or direction of change that
89 arises within sequences of measurements.

90

91 As an example, a simple cultural evolution study might involve one participant completing a
92 task, set by the experimenter, in front of an observer. Upon completion of the task, the first
93 participant's performance is evaluated, and the observer takes over the role of task
94 completion, with a new participant arriving to take the role of observer. This would generally
95 continue for a pre-specified number of iterations, which together would represent a single
96 replicate within the overall experimental design. In an example such as this, any changes in
97 the task scores would likely represent a key measure of interest.

98

99 It is not within the scope of the current review to provide an exhaustive catalogue of such
100 studies (and indeed more comprehensive reviews of the literature can be found elsewhere,
101 [2,3]). We instead intend to provide an overview of dominant approaches, using illustrative
102 examples of particular studies where relevant, with particular focus on those that permit
103 insights into the role of innovation.

104

105 In terms of the methods of structuring the multi-participant replicates in cultural evolution
106 experiments, some common approaches to this have been described in the previous literature.
107 Mesoudi [4] distinguished three main approaches, labelling these as the *transmission chain*
108 method, the *constant-group* method, and the *replacement* method. In transmission chain
109 studies, participants take part in the experimental task one at a time, in strict succession,
110 receiving information only from their immediate predecessor. In contrast, in studies using the
111 constant group method, all members of a replicate take part simultaneously, so group
112 membership is fixed and there is no addition of naïve participants. Although in all constant

113 group studies it is possible to learn from any other member of the group, a further distinction
114 can be drawn between one type of design, in which the exchange of information is
115 unrestricted (sometimes referred to in the literature as “open diffusion”, e.g. [5]), and those
116 where information exchange is under control within the experiment. Finally, the replacement
117 method incorporates elements of both transmission and constant group methods: in these
118 studies, a small group of participants complete the experimental task simultaneously (as with
119 constant groups), but experienced members of the group are replaced at regular intervals by
120 naïve newcomers, by way of simulating generational succession within a population. For this
121 reason, such approaches are also sometimes referred to as microsocieties [6,7]. Using this
122 method it is therefore possible to ensure complete turnover of group membership whilst
123 retaining some flexibility over whom participants can learn from.

124

125 Within the current review we intend to restrict our discussion to those studies which
126 incorporate generational turnover as part of the design (i.e. including transmission chains and
127 replacement microsocieties, but excluding studies using the constant group method).

128 Attributing changes that occur within constant groups to the process of cultural evolution
129 (characterized as a Darwinian process consisting of the selective retention of favourable
130 socially learnt cultural variants as well as a variety of non-selective processes such as drift,
131 migration, and invention, e.g. [8,9,10]) is relatively problematic, since individual learning
132 processes (particularly feedback from trial and error) will typically tend to result in
133 directional changes in behaviour over time. This makes it difficult to determine the extent to
134 which any such changes have occurred as a consequence of cultural evolution or merely the
135 effects of iterative individual learning. Such designs can nonetheless be extremely valuable
136 for certain research questions within this field (e.g. for comparing the effects of different
137 group sizes, to understand the additive effects of social information on individual learning,

138 e.g. [11] or for exploring how cultural traditions, once formed, are actually spread through
139 populations, e.g. [5]). However, for the purposes of identifying innovations (see next section)
140 we feel that transmission chain and replacement designs allow potential heuristics for doing
141 this, which are less readily interpretable in the context of constant group approaches.

142

143

144 2. Identifying innovation in studies of cultural evolution

145

146 As noted previously, the experimental studies of cultural evolution which we review here
147 were not designed explicitly for the purpose of investigating innovation, so the researchers
148 who have carried out these studies have typically not provided their own definitions of what
149 constitutes an innovation in the context of particular studies. In order to re-interpret the
150 results of those studies we need to define what we consider an innovation in a manner that we
151 can apply to all studies.

152

153 We therefore propose to take a pragmatic approach to identifying innovation in studies of
154 cultural evolution by taking the perspective of the outcome rather than the intention. We can
155 infer innovations indirectly by considering similarity measures which have been used as a
156 proxy for transmission fidelity, i.e. only cultural variants which differ sufficiently from
157 already existing variants (i.e. possessing a low similarity score) are considered innovations.

158

159 2.1 Measures of similarity

160

161 The usefulness of our definition of innovation rests on the ability to define the degree of
162 similarity between different cultural variants. Similarity has been explicitly quantified in a

163 number of studies, using a range of different methods. There are several reasons why
164 researchers have employed such measures as a dependent variable in their designs. In some
165 cases the motivation has been to determine whether material becomes more learnable with
166 transmission, as evidenced by decreasing error rates (and increasing similarity) over
167 generations [12,13,14]. In other cases, similarity estimates have been used to determine
168 whether performance improvements over generations are associated with a pattern of descent
169 with modification, indicative of cultural evolutionary processes [15]. Such measures can also
170 be used to establish whether separate lineages of variants are distinguishable from one
171 another, in a manner characteristic of distinctive cultural traditions [7,15,16,17].

172

173 The precise method used to evaluate similarity between variants is determined largely by the
174 nature of the behaviours in question. For example, in studies using artificial language
175 learning tasks, where the cultural variants being studied are sequences of linguistic symbols
176 (i.e. words, or sequences of words), Levenshtein edit distance has been used [12,13]. This
177 metric calculates how similar one string of characters is to another by counting the minimum
178 number of characters that must be substituted, inserted or deleted to transform one string into
179 another, normalised by the length of the longer string. Other studies have used subjective
180 judgements of similarity as assigned by naïve raters, by simply asking them to compare two
181 items and indicate how closely they resemble one another; this has the advantage of validity
182 as a direct measure of human perception of resemblance, but has the drawback of being
183 opaque in relation to the source of similarity in terms of which features are shared [7,15,16].
184 Verhoef et al. [14] used a similarity metric based on the acoustic physical properties of an
185 auditory signal, but derived the weightings assigned to these properties from perceptual
186 ratings obtained in a separate pilot study, thus using an objective measure with accompanying
187 assurance of subjective validity. In other studies similarity between variants, although not

188 explicitly part of the research design, can sometimes be inferred from other measures used to
189 track retention of particular features of interest (often those that were present in stimulus
190 material presented to the first generation of participants), by considering the number, or
191 proportion, of shared features (e.g. in “serial reproduction” studies, e.g. [18,19,20,21]). Based
192 on the used measure of similarity it seems plausible to quantify the rates of cultural change
193 and therefore the rates of innovation in different experimental studies.

194

195 2.2 Sources of innovation

196

197 Our definition of innovation does not distinguish between different sources of innovations.
198 In the modelling literature innovations are generally regarded to be a potential outcome of
199 individual learning [22] or of erroneous cultural transmission [8,23], the latter being
200 commonly referred to as mutation. While the exact characterization of an innovation varies
201 between approaches (e.g. sometimes defined as novel to the individual, and in other cases
202 defined more narrowly as novel to the population), their function is very similar: innovations
203 induce the possibility of cultural change into the considered system.

204

205 Within studies of cultural evolution therefore, “innovations” may similarly arise as a
206 consequence of transmission error, or individual learning (involving intentional invention or
207 modification on the part of the participant). However in current studies of cultural evolution it
208 will generally be difficult to distinguish between transmission error and individual learning
209 based on the available data (e.g. the sequence of cultural variants produced in a transmission
210 chain). Potential inferences about the source of innovations will depend on the chosen
211 experimental design (discussed in the next section).

212

213 This leads us to define innovativeness as a continuum, rather than a dichotomy, with faithful
214 social transmission and innovation considered as opposite ends of a spectrum of possibilities
215 representing a balance between the two. So for our purposes someone who intended to copy,
216 but who failed and produced something very different from anything to which they had been
217 exposed, would be defined as having innovated. In contrast, an individual who independently
218 conceived of a solution that was highly similar to another solution potentially available to
219 them via social learning would be defined as not having innovated. These simplifying
220 assumptions allow us to operationalise innovation in a way that makes it possible for us to
221 identify it from experimental studies of cultural evolution.

222

223 3. Sources and effects of innovation across study designs

224

225 In this section we consider how the design of cultural evolution experiments influences what
226 we can infer about the role of innovations in generating directional cultural change. We
227 review experimental studies of cultural evolution, to consider first of all what is the ostensive
228 goal from the perspective of the participant, i.e. how has “success” on the task been framed
229 by the experimenter? This aspect of the design has important implications for the source of
230 innovations, and whether these arise primarily as a consequence of imperfect reproduction
231 (i.e. learning errors), or learning errors plus intentional modification on the part of the
232 participant. Secondly, we also consider the measure of interest used by the researcher to
233 quantify the predicted cultural change. Depending on the type of change that is being tracked
234 over transmission, the effects of innovation may be either highly predictable, or relatively
235 unpredictable, in terms of the likelihood of shifting behaviour in the predicted direction of
236 change.

237

238 3.1 Task aims and incentives

239

240 In this section we discuss two broad categories of cultural studies which differ in terms of the
241 goal as presented to the participants. Specifically, we distinguish between studies requiring
242 accurate reproduction (denoted reproduction goal studies) and studies involving evaluation of
243 performance on a specified task (denoted performance goal studies).

244

245 **Reproduction goal studies.** In many studies of cultural evolution, the goal of the participant
246 is simply to reproduce material that is presented to them as accurately as possible. Studies of
247 this type date far back in the scientific literature, including notably Bartlett’s experiments
248 using the “method of serial reproduction” [18]. These studies typically involve a transmission
249 chain design, within which the first participant is presented with some original stimulus
250 material, and subsequent participants are presented with the reproduction produced by their
251 predecessor in the chain. More recent examples of this type of design include Mesoudi,
252 Whiten and Dunbar’s [19] study of the transmission of written narratives, Tan and Fay’s [20]
253 study of the transmission of spoken narratives, and Tamariz and Kirby’s [24] study of the
254 transmission of meaningless drawings.

255

256 There are also other research designs which frame the object of the task as being accurate
257 reproduction, but which assess this in slightly different ways involving probing completeness
258 of knowledge of the stimulus material, rather than rote reproduction. For example, in several
259 recent studies of the cultural evolution of languages (e.g. [12]), participants have been
260 exposed to a stimulus set of signal-meaning pairings, with their knowledge of this artificial
261 language assessed through their recall of the appropriate signal to attach to a particular
262 meaning (with the participant’s pairings then used as stimulus material for their successor).

263

264 In these studies, perfect reproduction effectively constitutes maximum success on the task, so
265 all participants should be aiming to copy their stimulus material as accurately as possible. In
266 such contexts, the only “innovations” that arise do so as a consequence of errors in social
267 learning, rather than individual learning. Furthermore, the “adaptation” that occurs represents
268 adaptation only to the cognition of the learners. A chain that culminated in the transmission
269 of material which was perfectly reproducible, without error, could in this sense be envisaged
270 as having reached a stable equilibrium in relation to this adaptive force (see experiment 1 in
271 [12] for an illuminating example which comes close to such a state).

272

273 **Performance goal studies.** In other studies aiming to document the effects of cultural
274 evolution, participants are not explicitly instructed to copy the material they are presented
275 with. Typically in such research designs there is some other goal (sometimes implicit, but
276 often relatively explicit in the participants’ instructions) related to a particular task, the
277 achievement of which corresponds to successful performance. In such studies, which may
278 involve a replacement micro-society or transmission chain design, information about the
279 efforts of other participants is simply available as a potential source of evidence about how
280 the task can be approached. Examples of this type of study include Caldwell and Millen’s
281 [15] study of paper aeroplane and spaghetti tower building in replacement micro-societies, in
282 which the participants’ objective was to maximise the flight distance of their plane or the
283 height of their tower. In this study task success was highly explicit, and no social information
284 was provided to the first participant in each micro-society. In other studies task success has
285 sometimes been more implicit, and these have generally involved an initial demonstration by
286 the experimenter for the first generation of participants. For example, in Flynn and Whiten’s
287 [25] study of three and five year old children, a demonstration was provided for the first

288 participant of each transmission chain, showing how beads could be extracted from the
289 experimental apparatus using a tool. The instructions to participants were simply that they
290 could “have a go” once it was their turn. Nonetheless, the objective of bead extraction must
291 have been apparent to the participants, many of whom were successful in achieving this goal
292 (including 50% of the five year olds in the control group, who had not even witnessed a
293 demonstration).

294

295 In contrast to studies in which the participant’s goal is accurate reproduction, innovations that
296 occur when the goal is task success are liable to include the effects of intentional invention
297 and modification as well as errors in social learning. Likewise, any adaptation occurs in
298 response to the demands of the task in question, as well as the learners’ general cognitive
299 biases.

300

301 3.2 Measures of adaptation

302

303 To our knowledge, in all experimental studies of cultural evolution, there is generally some
304 sort of expectation about the nature of the change that repeated transmission is liable to
305 generate. The different measures used, however, will be affected differently by innovations,
306 and in some cases, innovations arising from intentional modification are likely to affect
307 measures differently from those that arise from social learning errors. In the following section
308 we discuss three broad methods which have been used to measure adaptation in studies of
309 cultural evolution.

310

311 **Loss/distortion measures.** In many studies, the measure of interest simply involves tracking
312 the retention of source material which is presented to the first participant of a chain.

313 Examples of such measures include the number of propositions from the original stimulus
314 material which were accurately reproduced by participants in serial reproduction studies (e.g.
315 [19,20,21]). Although this type of measure is more commonly used in study designs
316 involving an explicit reproduction goal, measures tracking the retention of particular task
317 solutions are also sometimes used in studies which present participants with a task success
318 goal. For example, Flynn and Whiten's [25] study, mentioned previously, involved the study
319 of transmission chains which had been seeded with one of two different methods of using the
320 tool and apparatus. The study tracked the longevity of these alternative techniques over
321 repeated transmission.

322

323 In studies which use relatively straightforward retention measures, such as those described
324 above, innovations (which are necessarily deviations from retention) will have predictable
325 effects on the overall direction of change, increasing distortion and loss of information in
326 typically irreversible ways. Furthermore, the effects will occur regardless of whether the
327 innovations arise from individual learning or errors in social transmission, since any changes
328 will result in dilution and/or distortion of the source.

329

330 **Task success measures.** In studies where task success is the goal of the participant, this same
331 task success measure may be used to track changes as a consequence of transmission.

332 Generally, in designs where the first generation of participants have no social information,
333 the expectation would be that task success would tend to increase with transmission,
334 indicative of cumulative culture (e.g. [15]; image generation in [26]). In other designs, where
335 the chain is seeded with a demonstration from a skilled expert (e.g. knot-tying in [26]; [27]),
336 the task success measure is used to assess resistance against loss under different conditions of
337 transmission. Alternatively, in some studies, the chain may be seeded with a response that is

338 intentionally extreme in its ineffectiveness, or degree of error. This has allowed researchers to
339 investigate the persistence of, and recovery from, initially disadvantageous responses. For
340 example, Flynn [28] and McGuigan & Graham [29] studied the loss of irrelevant actions
341 from children's actions on a puzzle box task, in chains which had been seeded with a
342 demonstration including both necessary and unnecessary actions.

343

344 The effect of innovations on measures of task success is likely to be much more
345 unpredictable, compared with the effect that these have on straightforward measures of
346 retention. Intuitively, we would expect that errors in social learning would tend to reduce task
347 success measures. If the participant is attempting to copy (rather than intentionally
348 innovating) then they have presumably concluded that they are unlikely to be able to improve
349 upon the solution which is available to them via social learning, and although fortuitous
350 learning errors are not impossible, they are probably relatively rare. In contrast, asocial
351 processes of intentional invention and modification must be largely responsible for the
352 increases in task success observed in experimental studies of cumulative culture, and as such
353 it can clearly have positive effects on these measures. However, since the effects of novel
354 variants are necessarily unpredictable, this is by no means guaranteed, and it is likely that
355 intentional modifications also reduce task success measures in many instances. In Section 4
356 we return to this issue, to examine particular studies which may provide insights into the
357 relationship between innovation rate and task success measures.

358

359 **Cultural attractor measures.** In a third category of studies, the measure of interest
360 represents a specific property of the transmitted behaviour, which is predicted to increase
361 with transmission as a consequence of this property rendering the material more learnable.
362 The property is therefore assumed to represent some sort of cultural attractor [30] whose

363 presence, or probability, will tend to increase relative to source material provided to the first
364 participant of a chain (in which the attractor would be normally be represented at statistical
365 chance level or below) or in which the degree of representation might be systematically
366 varied, e.g. [31]. Examples of studies using this kind of measure include artificial language
367 learning studies which predict increases in structural compositionality [12], predictability of
368 grammatical markers [32], or regularisation [31]. However, we would also include in this
369 category studies which seek evidence of the emergence of cognitive "priors" over repeated
370 transmission [33,34]. In these studies, participants attempted to infer a function [33] or
371 category membership hypothesis [34] from a set of exemplar data, with their selected
372 function or hypothesis being used to generate exemplar data for the next participant. Over
373 repeated transmission, the functions and hypotheses which increased in probability were
374 those which represented known human learning biases.

375

376 These studies have typically emphasised a goal of accurate reproduction for participants
377 assessed by probing their knowledge of the learned material. However, it is also possible to
378 measure these sorts of changes in studies framed in terms of task success (e.g. see [13], for an
379 example of a language evolution study using effective dyadic communication as the
380 participants' goal).

381

382 Considering these studies from the perspective of the effects of innovation illustrates an
383 important difference between this type of measure of interest, and those involving simply loss
384 and/or distortion of source material. In studies looking for the emergence of cultural
385 attractors, it is perfectly possible for errors in transmission to result in changes which move in
386 in the opposite direction to the prediction. As noted previously, in studies documenting
387 degradation of source material, any kind of loss or distortion effectively generates change in

388 the predicted direction. Nonetheless, in studies measuring the presence of presumed cultural
389 attractors, it is still quite likely that increased error rates will tend to increase the cultural
390 change in the direction of the proposed attractor, since it is assumed to be the result of some
391 kind of cognitive bias.

392

393 Figure 1 provides an overview of the different categories of cultural evolution experiments
394 we have outlined here, i.e. in relation to the participant's goal, and the researcher's measure
395 of adaptation. The probable sources of innovation are specified for each, as well as their
396 likely effects on the measure of interest. It is worth noting that studies may actually report
397 multiple measures of adaptation as defined here. Depending upon the design it is possible in
398 principle to simultaneously track the retention of features from source material, the actual
399 performance in a given task, and the transitioning structural properties of the behaviour being
400 transmitted itself.

401

402

403 4. Rates of innovation and rates of change and adaptation

404

405 Lastly we turn to the existing evidence for the effects of innovations on the measures of
406 cultural adaptation. Although the intuitive assumption might be that higher rates of
407 innovation are likely to generate faster rates of cultural change and adaptation, this is not
408 necessarily the case. The direction of modifications arising from innovations may not be
409 consistent, potentially resulting in limited overall change despite low similarity between
410 traits. In this section we consider examples of studies of cultural evolution within which
411 differing rates of retention have been identified across experimental conditions, with a view
412 to assessing the validity of our expectations about the varying effects of innovation across

413 different study types (as outlined in Figure 1, and the previous section). We finish by
414 considering evidence from theoretical models, which serves to highlight important
415 distinctions between the structure of the models and the simplifying constraints within much
416 of the existing experimental work, which impact upon the role of innovation in adaptive
417 change.

418

419 4.1 Studies measuring loss or distortion of a source

420

421 For studies in which participants are presented with a goal of accurate reproduction, variation
422 in retention rates may be found as a result of the ease or difficulty with which this can be
423 achieved. In studies using the serial reproduction method for example, alternative methods of
424 presenting the stimulus material may facilitate more accurate duplication. Tan and Fay [20],
425 for example, compared the transmission of short narratives under two different conditions. In
426 one condition, participants listened to an audio recording of their predecessors' narration
427 (from recall) of a passage, and then produced their own recording from memory for their
428 successor. In the other condition, participants actually met and interacted with their
429 predecessor in the chain, receiving the account in person in the context of a conversation.
430 Recall was found to be better in the interactive condition. Similarly, Eriksson and Coultas
431 [21] also identified differing retention rates across experimental conditions, finding that
432 narratives were transmitted with higher fidelity when participants received the story from two
433 different individuals, compared with receiving a single individual's reproduction twice.

434

435 The effects that these different retention rates have on the measures of cultural adaptation
436 used in the studies is very much in line with the predictions detailed in Figure 1, with these
437 studies finding that lower retention generates more rapid loss of detail. This in itself is

438 unsurprising given that between-generation similarity and overall cultural change are
439 effectively being inferred from the same data (i.e. the presence or absence of details from the
440 source material). However, given that we can be relatively confident about the source of
441 innovations in these studies (copying error, as opposed to intentional innovation) these
442 studies also provide an insight into baseline levels of change that should be expected from
443 imperfect transmission alone. This information is useful from the point of view of identifying
444 the role of intentional innovation in other studies.

445

446 When participants are given a goal of success on a particular task, rather than a goal of
447 reproduction, it is possible to find variation in retention rates across experimental conditions
448 as a consequence of strategic choice as well as ease of reproduction. However, as detailed in
449 Section 3 and Figure 1, any such strategic shifts ought to have equally predictable effects on
450 measures of loss or distortion of source material.

451

452 One example of such an effect comes from Caldwell and Eve's [35] study of participants'
453 designs in a spaghetti tower building task. Participants were encouraged to build their towers
454 to be as tall as possible, in two conditions. In the control condition, participants were told
455 their reward payment was based on the final height of their tower. In the other
456 ("unpredictable payoff") condition, participants were told their tower would be subjected to
457 unspecified structural tests before being measured for payment, although in reality, no such
458 tests were carried out. The aim of the experiment was to track the influence of particular
459 tower designs which had been presented to the very first generation of participants, and to
460 determine whether the influence of the seeded designs would persist for longer under
461 conditions of uncertainty about payoffs for novel solutions (in line with a "copy when
462 uncertain" strategy, [36]). Members of transmission chains were shown photographs of the

463 towers produced by their two immediate predecessors, which they could choose to copy or
464 not, presumably based on their assessment of the likely utility of this information in relation
465 to the task goal. The overall prediction was supported, with towers in the unpredictable
466 reward condition showing higher between-generation similarity (as evaluated by number of
467 shared features), and evidence of residual similarity to the original seed towers in later
468 generations. This contrasted with the findings from the control condition, in which between-
469 generation similarity was lower, and there was no detectable influence of the seed designs in
470 later generations.

471

472 4.2 Studies measuring task success

473

474 Although the studies discussed above (Section 4.1) offer relatively unsurprising relationships
475 between rates of innovation and rates of change, this relationship is definitely appears to be
476 less straightforward in other study designs. In studies tracking measures of success on a
477 particular task presented to participants, variation in retention rates may again arise from
478 strategic shifts in the degree of reliance placed on social versus individual learning, but this
479 may not necessarily translate to different rates of adaptation. In one example of such a study,
480 Caldwell and Millen [7] aimed to build upon previous [15] work, which had identified
481 cumulative improvement in spaghetti tower building over generations of replacement
482 microsocieties, by incorporating an experimental manipulation designed to emphasise the
483 importance of tower stability as well as height. Similarly to [35], Participants in the stability-
484 emphasis condition were informed that their tower would be measured following a delay
485 during which structural resilience would be under threat. The resulting uncertainty about the
486 likely effectiveness of different designs appeared to generate a strategic shift towards greater
487 reliance on social information, with towers from this condition being rated as having higher

488 relative within-chain similarity, compared with those built by participants given a
489 straightforward height goal.

490

491 The critical question then is how the greater reluctance to innovate impacted on the goal
492 measure of tower height. Interestingly, participants in this condition did not appear to have
493 been placed at a disadvantage in terms of the height of their towers, which did not differ
494 significantly across conditions. And although evidence of cumulative improvement was
495 somewhat clearer in the condition favouring greater innovation (height emphasis only), there
496 was also evidence of height increases over generations in the stability emphasis condition, in
497 spite of the apparent conformity to particular design types.

498

499 Caldwell and Eve [35] followed this up using the seeded-chain design described previously,
500 which explored the persistence of particular designs across two experimental conditions
501 intended to correspond to the predictable and unpredictable contexts from [7]. As already
502 noted, the expectation regarding relative retention rates was supported, by examining the
503 retention rate of features from the seeded tower designs, but it was also possible to measure
504 task success in the shape of tower height. Consistent with the earlier [7] findings, there was
505 no clear difference between these two conditions in terms of success on the task. In contrast,
506 the specific design used to seed the chains (one of which was superior to the other) had a
507 clear effect on the height of the subsequent towers, common across both of the experimental
508 conditions.

509

510 Thus, in both cases, these strategic shifts in the balance between social and individual
511 learning have not been associated with an obvious advantage to greater innovation. This is
512 despite the fact that it must be differences in the likelihood of intentional innovation, rather

513 than the likelihood of error, which accounts for the differences between conditions. In
514 addition, it is worth noting that in both of these studies, the conditions exhibiting lower
515 innovation were ones which in reality needlessly constrained participants' choices
516 (particularly in [35], in which the task was simply framed differently across conditions, and
517 there was no real difference in the way the efficacy of designs was evaluated). In this context,
518 one might expect that there should be a clear advantage to participants in the conditions
519 which simply emphasised maximising height, without needing to consider trade-offs with
520 probable stability. However, bearing in mind that social learning is critical to the *retention* of
521 advantageous variants, this may explain why the greater willingness to explore alternatives
522 did not appear to generate benefits at group level, since this necessarily occurred at the
523 expense of the potential for retaining beneficial traits. Overall, these studies certainly provide
524 support for the expectation that innovation rates will not have a straightforward relationship
525 with measures of adaptation focussed on task success (Figure 1).

526

527 4.3 Studies measuring presence of a cultural attractor

528

529 In a recent study of the cultural evolution of structural simplicity, Kempe et al. [17] compared
530 transmission chains of children with adults, the participants' goal being to reproduce the
531 positioning of random dot patterns on a grid. The hypothesis was that patterns would simplify
532 more in the chains of children, as measured by the clustering of dots and algorithmic
533 complexity. The similarity between adjacent responses could also be assessed, based on the
534 percentage of dots correctly placed on the grid. Thus, it is possible to determine from the data
535 whether greater adaptation was associated with lower levels of similarity. Interestingly, in
536 spite of strong support for the hypothesis that simplification would be stronger in chains of
537 children, there was no difference between the two populations in the between-generation

538 error measures. This suggests that increased rate of simplification was not simply attributable
539 to the children making more errors, so in fact they must have made qualitatively different
540 errors, that were more likely to shift responses in the direction of greater structural simplicity.
541 So, although it is not possible to say from these data what effect an increase or decrease in
542 error rate might have had in relation to the rate of adaptation in either adults or children, this
543 clearly demonstrates that similar error rates do not necessarily dictate equivalent rates of
544 adaptation.

545

546 Currently, there appears to be very limited evidence of the effects of different rates of
547 retention on measures involving proposed cultural attractors. As noted in Section 3, there is
548 good reason to believe that higher innovation rates might be associated with more rapid
549 change in the direction of the cultural attractor. However, this remains to be established.
550 Furthermore, it is likely that, as with task success measures, relatively faithful retention may
551 be critical to preserving change in a particular direction, yielding a U-shaped relationship
552 between rates of innovation and adaptation. Further research could clarify the nature of this
553 relationship.

554

555 4.4 Insights from theoretical models

556

557 Theoretical work has suggested that there exists a trade-off between the amount of innovation
558 and the level of adaptation depending on the level of environmental stability. In contrast with
559 the experimental approaches, where the environment (the physical or cultural environment to
560 which the considered cultural trait becomes adapted) is typically held constant, mathematical
561 or computational models can manipulate this variable. Using this approach, it has been shown
562 that asocial and social learning are favoured by natural selection when temporal

563 environmental changes occur in short and long intervals, respectively [8,37,38,39,40,41]. In
564 other words, the faster the adaptive value of a cultural variant is changing the more
565 advantageous is the individual learning strategy. As individual learning is considered as the
566 innovation mechanism this result also points to the crucial relationship between
567 environmental stability and the amount of innovation/cultural variation that is needed to adapt
568 to those changing conditions. Innovations (in particular adaptive innovations) provide the
569 basis for social learning to be a successful evolutionary strategy even in changing
570 environments [23,42]. However, due to the possible adaptive and non-adaptive nature of
571 innovations, there exists an optimal balance between the rate of innovation (expressed by the
572 fraction of the population engaged in individual learning) and environmental uncertainty [42].
573 The more unstable the environment the higher the amount of variation needed to ensure
574 efficient adaptation. Naturally this relationship is greatly influenced by the specific social
575 learning strategy [41,43,44,45,46,47].

576

577 It is not obvious how to relate the insights from theoretical models directly to those generated
578 by the experimental studies, but consideration of the reasons for this difficulty highlights
579 constraints and assumptions within the experimental designs. Within the modelling literature,
580 innovations are generally viewed as a means of cultural change and, in particular, a means of
581 tracking environmental change. In contrast, in the experiments reviewed here, the
582 “environment” to which adaptation occurs is either the environment of the mind, or the task
583 plus the mind, and the studies document the process of approaching an equilibrium state,
584 from a starting point of either naivety, or from an experimentally induced non-equilibrium
585 state. However, taking this view, the varying effects of innovation rates on cultural adaptation
586 across different experimental designs can perhaps usefully be conceptualised as a
587 consequence of both the shape of the adaptive landscape, and the likelihood of innovations

588 climbing in the direction of local optima. Further research, of both a theoretical and
589 experimental nature, is needed to cross-validate specific conclusions.

590

591

592 5. Conclusions

593

594 In this paper we began by adopting a pragmatic definition of innovation (blind to the
595 motivations and intentions its creator) that would allow us to identify it within experimental
596 studies of cultural evolution. Nonetheless, we made the assumption that innovations arose
597 from two main sources in this respect, i.e. they were either the result of (unintentional) errors
598 in transmission, or intentional invention or modification on the part of the innovator. We then
599 used these assumptions and simplification to ask what can be learned from current
600 experimental studies about the process of innovation.

601

602 Based on the existing literature, only limited insights are possible. However, we can at least
603 compare rates of innovation between different studies, or between different conditions of a
604 single study, by considering measurements of similarity between variants. Studies which
605 present participants with a goal of accurate reproduction can in this respect provide us with
606 an indication of baseline levels of innovation that one should expect as a result of error alone
607 (although this will of course be highly dependent on the learnability of the material being
608 transmitted, so any generalisations to different contexts should be made with extreme
609 caution). Studies involving measures of task success can provide insights into the effect of
610 the balance between innovation and social learning on the rate of adaptation to the task
611 demands. Overall however, the existing literature does not yet provide a clear picture even in
612 relation to these issues. We believe that future experimental work would benefit from explicit

613 consideration of factors influencing innovation, and the effects that this has on the rate and
614 direction of cultural evolution.

615

616 As another extremely worthwhile avenue for future research, we believe that it should be
617 possible, at least in principle, to distinguish between intentional and unintentional innovation
618 in experimental studies similar to the ones we describe here. The ability to do so hinges on
619 differences in the degree of cultural variation produced by both sources of innovation. In
620 studies which present participants with a goal of achieving success on a particular task it
621 should be possible to quantify the expected amount of cultural variation due to error by
622 including a baseline condition requesting only accurate reproduction of previous solutions, in
623 place of task success. This would provide a benchmark to which observed variations could be
624 compared, with levels of similarity lower than the benchmark pointing to the presence of
625 processes of intentional innovation. We know of no study to date which has explicitly
626 compared the two types of task goal (although see [48] for a comparison between a
627 reproduction-goal transmission chain and real world data, which aims to draw a similar
628 inference). Such experiments would have the additional advantage of potentially revealing
629 which properties of cultural variants are most prone to modification as a consequence of
630 erroneous social learning. However, further research is clearly required in order to
631 substantiate these proposals.

632

633 In addition, we note that it is currently difficult to relate experimental work on this topic to
634 theoretical models which pose similar questions, due to differences in focus. We believe there
635 is a need for further research which attempts to bridge this gap in order to permit cross-
636 validation of results.

637

638 Additional Information

639

640 Authors' Contributions

641 All authors made substantial contributions to the initial conception of the article. CC drafted
642 the article and HC and AK revised it critically for important intellectual content. All authors
643 gave final approval of the version to be published.

644

645 Competing Interests

646 We have no competing interests.

647

648 Funding

649 C.A.C. was supported by the European Research Council (ERC) under the European Union's
650 Horizon 2020 research and innovation programme under grant agreement No. 648841
651 RATCHETCOG ERC-2014-CoG.

652

653 References

654

655 1. Tomasello, M. 1999 *The cultural origins of human cognition*. Cambridge, Mass.: Harvard
656 University Press.

657 2. Mesoudi, A. & Whiten, A. 2008 The multiple roles of cultural transmission experiments in
658 understanding human cultural evolution. *Phil. Trans. R. Soc. B* **363**, 3489-3501.

659 3. Whiten, A., Caldwell, C. A. & Mesoudi, A. 2016 Cultural diffusion in humans and other
660 animals. *Curr. Opin. Psychol.* **8**, 15–21.

661 4. Mesoudi, A. 2007 Using the methods of experimental social psychology to study cultural
662 evolution. *J. Soc. Evol. Cult. Psychol.* **12**, 35-58.

663 5. Flynn, E., & Whiten, A. 2010 Studying children’s social learning experimentally “in the
664 wild”. *Learn. Behav.* **38**, 284-296.

665 6. Baum, W. M., Richerson, P. J., Efferson, C. M. & Paciotti, B. M. 2004 Cultural evolution
666 in laboratory microsocieties including traditions of rule giving and rule following. *Evol.*
667 *Hum. Behav.* **25**, 305-326.

668 7. Caldwell C. A. & Millen, A. E. 2010 Conservatism in laboratory microsocieties:
669 unpredictable payoffs accentuate group-specific traditions. *Evol. Hum. Behav.* **31**, 123-
670 130.

671 8. Boyd, R. & Richerson, P. J. 1985 *Culture and the evolutionary process*. Chicago, IL:
672 Chicago University Press.

673 9. Cavalli-Sforza, L. L. & Feldman, M. W. 1981 *Cultural transmission and evolution: a*
674 *quantitative approach*. Princeton, NJ: Princeton University Press.

675 10. Laland, K. N. & Brown, G. R. 2011 *Sense and nonsense: evolutionary perspectives on*
676 *human behaviour, 2nd edition*. Oxford: Oxford University Press.

- 677 11. Derex, M., Beugin, M. P., Godelle, B., Raymond, M. 2013 Experimental evidence for the
678 influence of group size on cultural complexity. *Nature* **503**, 389-391.
- 679 12. Kirby, S., Cornish, H. & Smith, K. 2008 Cumulative cultural evolution in the laboratory:
680 an experimental approach to the origins of structure in human language. *PNAS* **10531**,
681 10681-10686.
- 682 13. Kirby, S., Tamariz, M., Cornish, H. & Smith, K. 2015 Compression and communication
683 in the cultural evolution of linguistic structure. *Cognition* **141**, 87-102.
- 684 14. Verhoef, T., Kirby, S. & de Boer, B. 2014 Emergence of combinatorial structure and
685 economy through iterated learning with continuous acoustic signals. *J. Phon.* **43**, 57-68.
- 686 15. Caldwell, C. A. & Millen, A. E. 2008 Experimental models for testing hypotheses about
687 cumulative cultural evolution. *Evol. Hum. Behav.* **29**, 65-71.
- 688 16. Matthews, C., Roberts, G. & Caldwell, C. A. 2012 Opportunity to assimilate and pressure
689 to discriminate can generate cultural divergence in the laboratory. *Evol. Hum. Behav.*
690 **33**, 759-770.
- 691 17. Kempe, V., Gauvrit, N., Forsyth, D. 2015 Structure emerges faster during cultural
692 transmission in children than in adults. *Cognition* **136**, 247-254.
- 693 18. Bartlett, F. C. 1932 *Remembering: a study in experimental and social psychology*.
694 Cambridge: Cambridge University Press.
- 695 19. Mesoudi, A., Whiten, A. & Dunbar, R. I. M. 2006 A bias for social information in human
696 cultural transmission. *Brit. J. Psychol.* **97**, 405-423.
- 697 20. Tan, R. & Fay, N. 2011 Cultural transmission in the laboratory: agent interaction
698 improves the intergenerational transfer of information. *Evol. Hum. Behav.* **32**, 399-406.
- 699 21. Eriksson, K. & Coultas, J. C. 2012 The advantage of multiple cultural parents in the
700 cultural transmission of stories. *Evol. Hum. Behav.* **33**, 251-259.

- 701 22. Rendell, L., Boyd, R., Cownden, D., Enquist, M., Eriksson, K., Feldman, M.W., Fogarty,
702 L., Ghirlanda, S., Lillicrap, T. P. & Laland, K. N. 2010 Why copy others? Insights from
703 the social learning strategies tournament. *Science* **328**, 208-213.
- 704 23. Bentley, R. A., Hahn, M. W. & Shennan, S. J. 2004 Random drift and cultural change.
705 *Proceedings of the Royal Society B* **271**, 1443-1450.
- 706 24. Tamariz, M. & Kirby, S. 2015 Culture: copying, compression, and conventionality. *Cogn.*
707 *Sci.* **39**, 171-183.
- 708 25. Flynn, E. & Whiten, A. 2008 Cultural transmission of tool use in young children: a
709 diffusion chain study. *Soc. Dev.* **17**, 699-718.
- 710 26. Muthukrishna, M., Shulman, B. W., Vasilescu, V. & Henrich, J. 2014 Sociality influences
711 cultural complexity. *Proc. R. Soc. B.* **281**, 20132511.
- 712 27. Morgan, T. J., Uomini, N. T., Rendell, L. E., Chouinard-Thuly, L., Street, S. E., Lewis, H.
713 M., Cross, C. P., Evans, C., Kearney, R., de la Torre, I., et al. 2015 Experimental
714 evidence for the co-evolution of hominin tool-making teaching and language. *Nat.*
715 *Commun.* **6**, 6029.
- 716 28. Flynn, E. 2008 Investigating children as cultural magnets: do young children transmit
717 redundant information along diffusion chains? *Philos. Trans. R. Soc. B.* **363**, 3541-
718 3551.
- 719 29. McGuigan, N. & Graham, M. 2010 Cultural transmission of irrelevant tool actions in
720 diffusion chains of 3- and 5-year-old children. *Eur. J. Dev. Psychol.* **7**, 561-577.
- 721 30. Sperber, D. 1996 *Explaining culture: a naturalistic approach*. Oxford: Blackwell.
- 722 31. Reali, F., Griffiths, T. L. 2009 The evolution of frequency distributions: relating
723 regularization to inductive biases through iterated learning. *Cognition* **111**, 317-328.
- 724 32. Smith, K. & Wonnacott, E. 2010 Eliminating unpredictable variation through iterated
725 learning. *Cognition* **116**, 444-449.

- 726 33. Kalish, M. L., Griffiths, T. L. & Lewandowsky, S. 2007 Iterated learning:
727 Intergenerational knowledge transmission reveals inductive biases. *Psychon. B. Rev.*
728 **14**, 288-294.
- 729 34. Griffiths, T. L., Christian, B. R. & Kalish, M. L. 2008 Using category structures to test
730 iterated learning as a method for identifying inductive biases. *Cogn. Sci.* **32**, 68-107.
- 731 35. Caldwell, C. A. & Eve, R. M. 2014 Persistence of contrasting traditions in cultural
732 evolution: unpredictable payoffs generate slower rates of cultural change. *PLoS One* **9**,
733 e99708.
- 734 36. Laland, K. N. 2004 Social learning strategies. *Learn. Behav.* **32**, 4–14.
- 735 37. Aoki, K., Wakano, J. Y. & Feldman, M. W. 2005 The emergence of social learning in a
736 temporally changing environment: a theoretical model. *Curr. Anthropol.* **46**, 334-340.
- 737 38. Aoki, K. & Feldman, M. W. 2014 Evolution of learning strategies in temporally and
738 spatially variable environments: a review of theory. *Theor. Popul. Biol.* **91**, 3–19.
- 739 39. Boyd, R. & Richerson, P. J. 1988 An evolutionary model of social learning: the effects of
740 spatial and temporal variation. In *Social Learning: Psychological and Biological*
741 *Perspectives* (eds T. Zentall & B.G. Galef), pp 29-48. Hillsdale, NJ: Lawrence Erlbaum
742 Associates.
- 743 40. Feldman, M. W., Aoki, K. & Kumm, J. 1996 Individual versus social learning:
744 evolutionary analysis in a fluctuating environment. *Anthropol. Sci.* **104**, 209–213.
- 745 41. Rogers, A. 1988 Does biology constrain culture? *Am. Anthropol.* **90**, 819–813.
- 746 42. Kandler, A. & Laland, K. N. 2013 Tradeoffs between the strength of conformity and
747 number of conformists in variable environments. *J. Theor. Biol.* **332**, 191-202.
- 748 43. Henrich, J., Boyd, R. 1998 The evolution of conformist transmission and the emergence
749 of between-group differences. *Evol. Hum. Behav.* **19**, 215–241.

- 750 44. Nakahashi, W. 2007 The evolution of conformist transmission in social learning when the
751 environment changes periodically. *Theor. Popul. Biol.* **72**, 52–66.
- 752 45. Wakano, J. & Aoki, K. 2007 Do social learning and conformist bias coevolve? Henrich
753 and Boyd revisited. *Theor. Popul. Biol.* **72**, 504–512.
- 754 46. Kendal, J. R., Giraldeau, L-A. & Laland, K. N. 2009 The evolution of social learning
755 rules: payoff-biased and frequency-dependent biased transmission. *J. Theor. Biol.* **2602**,
756 210-219.
- 757 47. Nakahashi, W., Wakano, J. Y. & Henrich, J. 2012 Adaptive social learning strategies in
758 temporally and spatially varying environments: how temporal vs. spatial variation,
759 number of cultural traits, and costs of learning influence the evolution of conformist-
760 biased transmission, payoff-biased transmission, and individual learning. *Hum. Nat.*
761 **234**, 386–418.
- 762 48. Kempe, M., Lycett, S. & Mesoudi, A. 2012 An experimental test of the accumulated
763 copying error model of cultural mutation for Acheulean handaxe size. *PLoS One* **7**,
764 e48333.
- 765

766 Figure and table captions

767

768 Figure 1. Sources of innovation in experimental studies of cultural evolution, and their likely
769 effects on measures of adaptation. See Section 4 for examples of studies within each of the
770 categories, based on the participant's goal and the researcher's measure of adaptation.

771 Upwards arrows indicate effects expected to promote the type of change being measured, and
772 downward arrows indicate effects expected to inhibit such changes. Large arrows indicate the
773 expected dominant force of change, and the presence of an additional smaller arrow indicates
774 the possibility of innovations also influencing the measure of adaptation in the opposite
775 direction to the expected dominant effect.

776

777

		Goal of participant				
		Accurate social learning		Task success		
Measure of adaptation	Loss and/or distortion of source material	Source of innovation:	Social learning error only	Source of innovation:	Social learning error	Intentional invention/modification
		Effect on measure:	↑	Effect on measure:	↑	↑
	Presumed cultural attractor arising from human cognitive biases	Source of innovation:	Social learning error only	Source of innovation:	Social learning error	Intentional invention/modification
	Effect on measure:	↑↓	Effect on measure:	↑↓	↑↓	
	Task success	N/A		Source of innovation:	Social learning error	Intentional invention/modification
				Effect on measure:	↓↑	↑↓