



Review

Cite this article: Caldwell CA, Cornish H, Kandler A. 2016 Identifying innovation in laboratory studies of cultural evolution: rates of retention and measures of adaptation. *Phil. Trans. R. Soc. B* **371**: 20150193. <http://dx.doi.org/10.1098/rstb.2015.0193>

Accepted: 9 December 2015

One contribution of 15 to a theme issue 'Innovation in animals and humans: understanding the origins and development of novel and creative behaviour'.

Subject Areas:

behaviour, cognition, evolution

Keywords:

cultural evolution, iterated learning, microsociey, social learning, transmission chain

Author for correspondence:

Christine A. Caldwell
e-mail: c.a.caldwell@stir.ac.uk

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

Christine A. Caldwell¹, Hannah Cornish¹ and Anne Kandler²

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

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

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

1. Introduction

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

Following these definitions, it is clear that understanding innovation is fundamental to understanding cultural change. It is widely acknowledged that both innovation and social learning are the two cornerstones of cultural evolution [1]. While faithful social learning (i.e. social learning without any source of error) operates to maintain cultural traditions, on its own it will produce only cultural stasis over time. Biased social learning can generate a limited degree of cultural change in terms of the proportion of the population exhibiting a particular trait, in the direction of a stable end point. However, innovation is the primary driver of cultural change. Understanding the contexts that promote innovation, and the effect this has on population-level shifts in behaviour, is therefore essential to understanding phenomena as diverse as developments in science and technology, the rise and fall of fads and fashions, and shifting societal trends.

There are now numerous experiments reported in the literature that purport to capture aspects of cultural evolution under laboratory conditions. Potentially, these should offer fertile ground for helping us understand the catalysts and consequences of innovation within populations of learners. To the best of our knowledge, none of these studies has been designed with the explicit intention of investigating innovation, as they are more concerned with documenting overall patterns of change, rather than identifying particular individuals, or particular individual decisions, as the source of such change. However, some studies do provide an insight into factors affecting rates of innovation. Furthermore, in studies

that permit inferences about variation in innovation rate, it is also possible to consider the effect this has on the measures of directional cultural change used by the researcher. Although cultural change requires innovations, it does not necessarily follow that high innovation rates generate pronounced cultural change, aggregated over multiple learners. Depending on the circumstances under consideration, innovations may not necessarily modify cultural traits in consistent directions, generating limited change at the group level. In this paper, we focus on experimental research on cultural evolution with the aim to review what we can infer about the role of innovation in these studies.

(a) Experimental studies of cultural evolution

Although experimental studies of cultural evolution may take a variety of forms, we believe that all designs share certain unifying features that are worth outlining here. First, in contrast to more typical psychological experiments that concern how a single individual performs on a task, or sometimes how one individual learns from another, in studies of cultural evolution a single replicate within an experiment consists of multiple (three or more) participants. In this way, these designs capture the *repeated* occurrences of social learning involved in cultural change, as opposed to one-off cases of individual learning or social learning in general. Second, within each replicate, participants have some form of access to information about the solutions or responses of other members of the same replicate. The exact nature of the information available may vary, but can include direct observation, verbal report or stored information about solutions or responses presented remotely, i.e. in the absence of their progenitor. Finally, all studies involve a measure that is repeated successively, the overall aim being to describe the nature and/or direction of change that arises within sequences of measurements.

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

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

In terms of the methods of structuring the multi-participant replicates in cultural evolution experiments, some common approaches have been described in the previous literature. Mesoudi [4] distinguished three main approaches, labelling these as the *transmission chain* method, the *constant group* method and the *replacement* method. In transmission chain studies, participants take part in the experimental task one at a time, in strict succession, receiving information only from their immediate predecessor. In contrast, in studies using the constant group method all members of a replicate

take part simultaneously, so group membership is fixed, and there is no addition of naive participants. Although in all constant group studies it is possible to learn from any other member of the group, a further distinction can be drawn between one type of design, in which the exchange of information is unrestricted (sometimes referred to in the literature as 'open diffusion' [5]), and those where information exchange is under control within the experiment. Finally, the replacement method incorporates elements of both transmission and constant group methods: in these studies, a small group of participants complete the experimental task simultaneously (as with constant groups), but experienced members of the group are replaced at regular intervals by naive newcomers, by way of simulating generational succession within a population. For this reason, such approaches are also sometimes referred to as microsocieties [6,7]. Using this method, it is therefore possible to ensure complete turnover of group membership while retaining some flexibility over whom participants can learn from.

Within the current review, we intend to restrict our discussion to those studies that incorporate generational turnover as part of the design (i.e. including transmission chains and replacement microsocieties, but excluding studies using the constant group method). Attributing changes that occur within constant groups to the process of cultural evolution (characterized as a Darwinian process consisting of the selective retention of favourable socially learnt cultural variants as well as a variety of non-selective processes such as drift, migration and invention [8–10]) is relatively problematic, because individual learning processes (particularly feedback from trial and error) will typically tend to result in directional changes in behaviour over time. This makes it difficult to determine the extent to which any such changes have occurred as a consequence of cultural evolution or merely as an effect of iterative individual learning. Such designs can nonetheless be extremely valuable for certain research questions within this field (e.g. for comparing the effects of different group sizes, to understand the additive effects of social information on individual learning [11] or for exploring how cultural traditions, once formed, are actually spread through populations [5]). However, for the purposes of identifying innovations (see §2), we feel that transmission chain and replacement designs allow potential heuristics for doing this, which are less readily interpretable in the context of constant group approaches.

2. Identifying innovation in studies of cultural evolution

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

We therefore propose to take a pragmatic approach to identifying innovation in studies of cultural evolution by taking the perspective of the outcome rather than the intention. We can infer innovations indirectly by considering similarity measures that have been used as a proxy for transmission fidelity, i.e. only cultural variants that differ sufficiently from already

existing variants (i.e. possessing a low similarity score) are considered innovations.

(a) Measures of similarity

The usefulness of our definition of innovation rests on the ability to define the degree of similarity between different cultural variants. Similarity has been explicitly quantified in a number of studies, using a range of different methods. There are several reasons why researchers have employed such measures as a dependent variable in their designs. In some cases, the motivation has been to determine whether material becomes more learnable with transmission, as evidenced by decreasing error rates (and increasing similarity) over generations [12–14]. In other cases, similarity estimates have been used to determine whether performance improvements over generations are associated with a pattern of descent with modification, indicative of cultural evolutionary processes [15]. Such measures can also be used to establish whether separate lineages of variants are distinguishable from one another, in a manner characteristic of distinctive cultural traditions [7,15–17].

The precise method used to evaluate similarity between variants is determined largely by the nature of the behaviours in question. For example, in studies using artificial language learning tasks, where the cultural variants being studied are sequences of linguistic symbols (i.e. words or sequences of words), Levenshtein edit distance has been used [12,13]. This metric calculates how similar one string of characters is to another by counting the minimum number of characters that must be substituted, inserted or deleted to transform one string into another, normalized by the length of the longer string. Other studies have used subjective judgements of similarity as assigned by naive raters, by simply asking them to compare two items and indicate how closely they resemble one another; this has the advantage of validity as a direct measure of human perception of resemblance, but has the drawback of being opaque in relation to the source of similarity in terms of which features are shared [7,15,16]. Verhoef *et al.* [14] used a similarity metric based on the physical acoustic properties of an auditory signal, but derived the weightings assigned to these properties from perceptual ratings obtained in a separate pilot study, thus using an objective measure with accompanying assurance of subjective validity. In other studies, similarity between variants, although not explicitly part of the research design, can sometimes be inferred from other measures used to track retention of particular features of interest (often those that were present in stimulus material presented to the first generation of participants), by considering the number, or proportion, of shared features (e.g. in ‘serial reproduction’ studies [18–21]).

By using such measures of similarity, it seems plausible to quantify the degree of deviation from perfect copying, and therefore the rates of innovation in different experimental studies.

(b) Sources of innovation

Our definition of innovation does not distinguish between different sources of innovations. In the modelling literature, innovations are generally regarded to be a potential outcome of individual learning [22] or of erroneous cultural transmission [8,23], the latter being commonly referred to as mutation. While the exact characterization of an innovation varies between approaches (e.g. sometimes defined as novel to the individual, and in other cases defined more narrowly

as novel to the population), their function is very similar: innovations induce the possibility of cultural change into the considered system.

Within studies of cultural evolution, therefore, innovations may similarly arise as a consequence of transmission error, or individual learning (involving intentional invention or modification on the part of the participant). However, in current studies of cultural evolution, it will generally be difficult to distinguish between transmission error and individual learning based on the available data (e.g. the sequence of cultural variants produced in a transmission chain). Potential inferences about the source of innovations will depend on the chosen experimental design (discussed in §3).

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

3. Sources and effects of innovation across study designs

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

(a) Task aims and incentives

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

(i) Reproduction goal studies

In many studies of cultural evolution, the goal of the participant is simply to reproduce material that is presented to them as accurately as possible. Studies of this type date far back in

the scientific literature, including notably Bartlett's experiments using the 'method of serial reproduction' [18]. These studies typically involve a transmission chain design, within which the first participant is presented with some original stimulus material, and subsequent participants are presented with the reproduction produced by their predecessor in the chain. More recent examples of this type of design include Mesoudi *et al.*'s [19] study of the transmission of written narratives, Tan & Fay's [20] study of the transmission of spoken narratives and Tamariz & Kirby's [24] study of the transmission of meaningless drawings.

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

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

(ii) Performance goal studies

In other studies aiming to document the effects of cultural evolution, participants are not explicitly instructed to copy the material they are presented with. Typically, in such research designs, there is some other goal (sometimes implicit, but often relatively explicit in the participants' instructions) related to a particular task, the achievement of which corresponds to successful performance. In such studies, which may involve a replacement micro-society or transmission chain design, information about the efforts of other participants is simply available as a potential source of evidence about how the task can be approached. Examples of this type of study include Caldwell & Millen's [15] study of paper aeroplane and spaghetti tower building in replacement micro-societies, in which the participants' objective was to maximize the flight distance of their plane or the height of their tower. In this study, task success was highly explicit, and no social information was provided to the first participant in each micro-society. In other studies, task success has sometimes been more implicit, and these have generally involved an initial demonstration by the experimenter for the first generation of participants. For example, in Flynn & Whiten's [25] study of 3- and 5-year-old children, a demonstration was provided for the first participant of each transmission chain, showing how beads could be extracted from the experimental apparatus using a tool. The instructions to participants were simply that they could 'have

a go' once it was their turn. Nonetheless, the objective of bead extraction must have been apparent to the participants, many of whom were successful in achieving this goal (including 50% of the 5-year-olds in the control group, who had not even witnessed a demonstration).

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

(b) Measures of adaptation

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

(i) Loss/distortion measures

In many studies, the measure of interest simply involves tracking the retention of source material that is presented to the first participant of a chain. Examples of such measures include the number of propositions from the original stimulus material that were accurately reproduced by participants in serial reproduction studies [19–21]. Although this type of measure is more commonly used in study designs involving an explicit reproduction goal, measures tracking the retention of particular task solutions are also sometimes used in studies that present participants with a task success goal. For example, Flynn & Whiten's [25] study, mentioned previously, involved the study of transmission chains that had been seeded with one of two different methods of using the tool and apparatus. The study tracked the longevity of these alternative techniques over repeated transmission.

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

(ii) Task success measures

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

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

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

(iii) Cultural attractor measures

In a third category of studies, the measure of interest represents a specific property of the transmitted behaviour, which is predicted to increase with transmission as a consequence of this property rendering the material more learnable. The property is therefore assumed to represent some sort of cultural attractor [30] whose presence, or probability, will tend to increase relative to source material provided to the first participant of a chain (in which the attractor would be normally be represented at statistical chance level or below) or in which the degree of representation might be systematically varied [31]. Examples of studies using this kind of measure include artificial language learning studies that predict increases in structural compositionality [12], predictability of grammatical markers [32] or regularization [31]. However, we would also include in this category studies that seek evidence of the emergence of cognitive 'priors' over repeated transmission [33,34]. In these studies, participants attempted to infer a function [33] or category membership hypothesis [34] from a set of exemplar data, with their selected function or hypothesis being used to generate exemplar data for the next participant. Over repeated transmission, the functions and hypotheses that increased in probability were those that represented known human learning biases.

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

Considering these studies from the perspective of the effects of innovation illustrates an important difference between this

type of measure of interest, and those involving simply loss and/or distortion of source material. In studies looking for the emergence of cultural attractors, it is perfectly possible for errors in transmission to result in changes that move in the opposite direction to the prediction. As noted previously, in studies documenting degradation of source material, any kind of loss or distortion effectively generates change in the predicted direction. Nonetheless, in studies measuring the presence of presumed cultural attractors, it is still quite likely that increased error rates will tend to increase the cultural change in the direction of the proposed attractor, because it is assumed to be the result of some kind of cognitive bias.

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

4. Rates of innovation and rates of change and adaptation

Lastly, we turn to the existing evidence for the effects of innovations on the measures of cultural adaptation. Although the intuitive assumption might be that higher rates of innovation are likely to generate faster rates of cultural change and adaptation, this is not necessarily the case. The direction of modifications arising from innovations may not be consistent, potentially resulting in limited overall change despite low similarity between traits. In this section, we consider examples of studies of cultural evolution within which differing rates of retention have been identified across experimental conditions, with a view to assessing the validity of our expectations about the varying effects of innovation across different study types (as outlined in figure 1, and the previous section). We finish by considering evidence from theoretical models, which serves to highlight important distinctions between the structure of the models and the simplifying constraints within much of the existing experimental work that impact upon the role of innovation in adaptive change.

(a) Studies measuring loss or distortion of a source

For studies in which participants are presented with a goal of accurate reproduction, variation in retention rates may be found as a result of the ease or difficulty with which this can be achieved. In studies using the serial reproduction method, for example, alternative methods of presenting the stimulus material may facilitate more accurate duplication. Tan & Fay [20], for example, compared the transmission of short narratives under two different conditions. In one condition, participants listened to an audio recording of their predecessors' narration (from recall) of a passage, and then produced their own recording from memory for their successor. In the other condition, participants actually met and interacted with their predecessor in the chain, receiving the account in

		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:	↓↑	↑↓

Figure 1. Sources of innovation in experimental studies of cultural evolution, and their likely effects on measures of adaptation. See §4 for examples of studies within each of the categories, based on the participant's goal and the researcher's measure of adaptation. Upwards arrows indicate effects expected to promote the type of change being measured, and downward arrows indicate effects expected to inhibit such changes. Large arrows indicate the expected dominant force of change, and the presence of an additional smaller arrow indicates the possibility of innovations also influencing the measure of adaptation in the opposite direction to the expected dominant effect.

person in the context of a conversation. Recall was found to be better in the interactive condition. Similarly, Eriksson & Coultas [21] also identified differing retention rates across experimental conditions, finding that narratives were transmitted with higher fidelity when participants received the story from two different individuals, compared with receiving a single individual's reproduction twice.

The effects that these different retention rates have on the measures of cultural adaptation used in the studies is very much in line with the predictions detailed in figure 1, with these studies finding that lower retention generates more rapid loss of detail. This in itself is unsurprising given that between-generation similarity and overall cultural change are effectively being inferred from the same data (i.e. the presence or absence of details from the source material). However, given that we can be relatively confident about the source of innovations in these studies (copying error, as opposed to intentional innovation), these studies also provide an insight into baseline levels of change that should be expected from imperfect transmission alone. This information is useful from the point of view of identifying the role of intentional innovation in other studies.

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

One example of such an effect comes from Caldwell & Eve's [35] study of participants' designs in a spaghetti tower building task. Participants were encouraged to build their towers to be as tall as possible, in two conditions. In the control condition, participants were told their reward

payment was based on the final height of their tower. In the other ('unpredictable payoff') condition, participants were told their tower would be subjected to unspecified structural tests before being measured for payment, although in reality, no such tests were carried out. The aim of the experiment was to track the influence of particular tower designs that had been presented to the very first generation of participants, and to determine whether the influence of the seeded designs would persist for longer under conditions of uncertainty about payoffs for novel solutions (in line with a 'copy when uncertain' strategy [36]). Members of transmission chains were shown photographs of the towers produced by their two immediate predecessors, which they could choose to copy or not, presumably based on their assessment of the likely utility of this information in relation to the task goal. The overall prediction was supported, with towers in the unpredictable reward condition showing higher between-generation similarity (as evaluated by number of shared features), and evidence of residual similarity to the original seed towers in later generations. This contrasted with the findings from the control condition, in which between-generation similarity was lower, and there was no detectable influence of the seed designs in later generations.

(b) Studies measuring task success

Although the studies discussed above (§4a) offer relatively unsurprising relationships between rates of innovation and rates of change, this relationship definitely appears to be less straightforward in other study designs. In studies tracking measures of success on a particular task presented to participants, variation in retention rates may again arise from strategic shifts in the degree of reliance placed on social versus individual learning, but this may not necessarily

translate to different rates of adaptation. In one example of such a study, Caldwell & Millen [7] aimed to build upon previous [15] work, which had identified cumulative improvement in spaghetti tower building over generations of replacement microsocieties, by incorporating an experimental manipulation designed to emphasize the importance of tower stability as well as height. Similarly to [35], participants in the stability-emphasis condition were informed that their tower would be measured following a delay during which structural resilience would be under threat. The resulting uncertainty about the likely effectiveness of different designs appeared to generate a strategic shift towards greater reliance on social information, with towers from this condition being rated as having higher relative within-chain similarity, compared with those built by participants given a straightforward height goal.

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

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

Thus, in both cases, these strategic shifts in the balance between social and individual learning have not been associated with an obvious advantage to greater innovation. This is despite the fact that it must be differences in the likelihood of intentional innovation, rather than the likelihood of error, which account for the differences between conditions. In addition, it is worth noting that in both these studies, the conditions exhibiting lower innovation were ones that in reality needlessly constrained participants' choices (particularly in [35], in which the task was simply framed differently across conditions, and there was no real difference in the way the efficacy of designs was evaluated). In this context, one might expect that there should be a clear advantage to participants in the conditions that simply emphasized maximizing height, without needing to consider trade-offs with probable stability. However, bearing in mind that social learning is critical to the *retention* of advantageous variants, this may explain why the greater willingness to explore alternatives did not appear to generate benefits at group level, because this necessarily occurred at the expense of the potential for retaining beneficial traits. Overall, these studies certainly provide support for the expectation that innovation rates will not have a straightforward relationship with measures of adaptation focused on task success (figure 1).

(c) Studies measuring presence of a cultural attractor

In a recent study of the cultural evolution of structural simplicity, Kempe *et al.* [17] compared transmission chains of children with adults, the participants' goal being to reproduce the positioning of random dot patterns on a grid. The hypothesis was that patterns would simplify more in the chains of children, as measured by the clustering of dots and algorithmic complexity. The similarity between adjacent responses could also be assessed, based on the percentage of dots correctly placed on the grid. Thus, it is possible to determine from the data whether greater adaptation was associated with lower levels of similarity. Interestingly, in spite of strong support for the hypothesis that simplification would be stronger in chains of children, there was no difference between the two populations in the between-generation error measures. This suggests that increased rate of simplification was not simply attributable to the children making more errors, so in fact they must have made qualitatively different errors that were more likely to shift responses in the direction of greater structural simplicity. So, although it is not possible to say from these data what effect an increase or decrease in error rate might have had in relation to the rate of adaptation in either adults or children, this clearly demonstrates that similar error rates do not necessarily dictate equivalent rates of adaptation.

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

(d) Insights from theoretical models

Theoretical work has suggested that there exists a trade-off between the amount of innovation and the level of adaptation depending on the level of environmental stability. In contrast with the experimental approaches, where the environment (the physical or cultural environment to which the considered cultural trait becomes adapted) is typically held constant, mathematical or computational models can manipulate this freely. Using this approach, it has been shown that asocial and social learning are favoured by natural selection when temporal environmental changes occur in short and long intervals, respectively [8,37–41]. In other words, the faster the adaptive value of a cultural variant is changing, the more advantageous is the individual learning strategy. As individual learning is considered as the innovation mechanism, this result also points to the crucial relationship between environmental stability and the amount of innovation/cultural variation that is needed to adapt to those changing conditions. Innovations (in particular adaptive innovations) provide the basis for social learning to be a successful evolutionary strategy even in changing environments [23,42]. However, owing to the possible adaptive and non-adaptive nature of innovations, there exists an optimal balance between the rate of innovation (expressed by the fraction of the population engaged in individual learning) and environmental uncertainty [42]. The more unstable the environment, the higher the amount of

variation needed to ensure efficient adaptation. Naturally, this relationship is greatly influenced by the specific social learning strategy [41,43–47].

It is not obvious how to relate the insights from theoretical models directly to those generated by the experimental studies, but consideration of the reasons for this difficulty highlights constraints and assumptions within the experimental designs. Within the modelling literature, innovations are generally viewed as a means of cultural change and, in particular, a means of tracking environmental change. In contrast, in the experiments reviewed here, the ‘environment’ to which adaptation occurs is either the environment of the mind, or the task plus the mind, and the studies document the process of approaching an equilibrium state from a starting point of either naivety, or from an experimentally induced non-equilibrium state. However, taking this view, the varying effects of innovation rates on cultural adaptation across different experimental designs can perhaps usefully be conceptualized as a consequence of both the shape of the adaptive landscape, and the likelihood of innovations climbing in the direction of local optima. Further research, of both a theoretical and experimental nature, is needed to cross-validate specific conclusions.

5. Conclusion

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

Based on the existing literature, only limited insights are possible. However, we can at least compare rates of innovation between different studies, or between different conditions of a single study, by considering measurements of similarity between variants. Studies that present participants with a goal of accurate reproduction can in this respect provide us with an indication of baseline levels of innovation that one should expect as a result of error alone (although this will of course be highly dependent on the learnability of the material

being transmitted, so any generalizations to different contexts should be made with extreme caution). Studies involving measures of task success can provide insights into the effect of the balance between innovation and social learning on the rate of adaptation to the task demands. Overall however, the existing literature does not yet provide a clear picture even in relation to these issues. We believe that future experimental work would benefit from explicit consideration of factors influencing innovation, and the effects that these have on the rate and direction of cultural evolution.

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

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

Authors’ contributions. All authors made substantial contributions to the initial conception of the article. C.A.C. drafted the article and H.C. and A.K. revised it critically for important intellectual content. All authors gave final approval of the version to be published.

Competing interests. We have no competing interests.

Funding. C.A.C. is funded by an ERC Consolidator Grant.

References

1. Tomasello M. 1999 *The cultural origins of human cognition*. Cambridge, MA: Harvard University Press.
2. Mesoudi A, Whiten A. 2008 The multiple roles of cultural transmission experiments in understanding human cultural evolution. *Phil. Trans. R. Soc. B* **363**, 3489–3501. (doi:10.1098/rstb.2008.0129)
3. Whiten A, Caldwell CA, Mesoudi A. 2016 Cultural diffusion in humans and other animals. *Curr. Opin. Psychol.* **8**, 15–21. (doi:10.1016/j.copsyc.2015.09.002)
4. Mesoudi A. 2007 Using the methods of experimental social psychology to study cultural evolution. *J. Soc. Evol. Cult. Psychol.* **12**, 35–58. (doi:10.1037/h0099359)
5. Flynn E, Whiten A. 2010 Studying children’s social learning experimentally ‘in the wild’. *Learn. Behav.* **38**, 284–296. (doi:10.3758/LB.38.3.284)
6. Baum WM, Richerson PJ, Efferson CM, Paciotti BM. 2004 Cultural evolution in laboratory microsocieties including traditions of rule giving and rule following. *Evol. Hum. Behav.* **25**, 305–326. (doi:10.1016/j.evolhumbehav.2004.05.003)
7. Caldwell CA, Millen AE. 2010 Conservatism in laboratory microsocieties: unpredictable payoffs accentuate group-specific traditions. *Evol. Hum. Behav.* **31**, 123–130. (doi:10.1016/j.evolhumbehav.2009.08.002)
8. Boyd R, Richerson PJ. 1985 *Culture and the evolutionary process*. Chicago, IL: Chicago University Press.
9. Cavalli-Sforza LL, Feldman MW. 1981 *Cultural transmission and evolution: a quantitative approach*. Princeton, NJ: Princeton University Press.
10. Laland KN, Brown GR. 2011 *Sense and nonsense: evolutionary perspectives on human behaviour*, 2nd edn. Oxford, UK: Oxford University Press.
11. Derex M, Beugin MP, Godelle B, Raymond M. 2013 Experimental evidence for the influence of group

- size on cultural complexity. *Nature* **503**, 389–391. (doi:10.1038/nature12774)
12. Kirby S, Cornish H, Smith K. 2008 Cumulative cultural evolution in the laboratory: an experimental approach to the origins of structure in human language. *Proc. Natl Acad. Sci. USA* **105**, 10 681–10 686. (doi:10.1073/pnas.0707835105)
 13. Kirby S, Tamariz M, Cornish H, Smith K. 2015 Compression and communication in the cultural evolution of linguistic structure. *Cognition* **141**, 87–102. (doi:10.1016/j.cognition.2015.03.016)
 14. Verhoef T, Kirby S, de Boer B. 2014 Emergence of combinatorial structure and economy through iterated learning with continuous acoustic signals. *J. Phon.* **43**, 57–68. (doi:10.1016/j.wocn.2014.02.005)
 15. Caldwell CA, Millen AE. 2008 Experimental models for testing hypotheses about cumulative cultural evolution. *Evol. Hum. Behav.* **29**, 65–71. (doi:10.1016/j.evolhumbehav.2007.12.001)
 16. Matthews C, Roberts G, Caldwell CA. 2012 Opportunity to assimilate and pressure to discriminate can generate cultural divergence in the laboratory. *Evol. Hum. Behav.* **33**, 759–770. (doi:10.1016/j.evolhumbehav.2012.06.004)
 17. Kempe V, Gauvrit N, Forsyth D. 2015 Structure emerges faster during cultural transmission in children than in adults. *Cognition* **136**, 247–254. (doi:10.1016/j.cognition.2014.11.038)
 18. Bartlett FC. 1932 *Remembering: a study in experimental and social psychology*. Cambridge, UK: Cambridge University Press.
 19. Mesoudi A, Whiten A, Dunbar RIM. 2006 A bias for social information in human cultural transmission. *Brit. J. Psychol.* **97**, 405–423. (doi:10.1348/000712605X85871)
 20. Tan R, Fay N. 2011 Cultural transmission in the laboratory: agent interaction improves the intergenerational transfer of information. *Evol. Hum. Behav.* **32**, 399–406. (doi:10.1016/j.evolhumbehav.2011.01.001)
 21. Eriksson K, Coultas JC. 2012 The advantage of multiple cultural parents in the cultural transmission of stories. *Evol. Hum. Behav.* **33**, 251–259. (doi:10.1016/j.evolhumbehav.2011.10.002)
 22. Rendell L *et al.* 2010 Why copy others? Insights from the social learning strategies tournament. *Science* **328**, 208–213. (doi:10.1126/science.1184719)
 23. Bentley RA, Hahn MW, Shennan SJ. 2004 Random drift and cultural change. *Proc. R. Soc. Lond. B* **271**, 1443–1450. (doi:10.1098/rspb.2004.2746)
 24. Tamariz M, Kirby S. 2015 Culture: copying, compression, and conventionality. *Cogn. Sci.* **39**, 171–183. (doi:10.1111/cogs.12144)
 25. Flynn E, Whiten A. 2008 Cultural transmission of tool use in young children: a diffusion chain study. *Soc. Dev.* **17**, 699–718. (doi:10.1111/j.1467-9507.2007.00453.x)
 26. Muthukrishna M, Shulman BW, Vasilescu V, Henrich J. 2014 Sociality influences cultural complexity. *Proc. R. Soc. B* **281**, 20132511. (doi:10.1098/rspb.2013.2511)
 27. Morgan TJ *et al.* 2015 Experimental evidence for the co-evolution of hominin tool-making teaching and language. *Nat. Commun.* **6**, 6029. (doi:10.1038/ncomms7029)
 28. Flynn E. 2008 Investigating children as cultural magnets: do young children transmit redundant information along diffusion chains? *Phil. Trans. R. Soc. B* **363**, 3541–3551. (doi:10.1098/rstb.2008.0136)
 29. McGuigan N, Graham M. 2010 Cultural transmission of irrelevant tool actions in diffusion chains of 3- and 5-year-old children. *Eur. J. Dev. Psychol.* **7**, 561–577. (doi:10.1080/17405620902858125)
 30. Sperber D. 1996 *Explaining culture: a naturalistic approach*. Oxford, UK: Blackwell.
 31. Reali F, Griffiths TL. 2009 The evolution of frequency distributions: relating regularization to inductive biases through iterated learning. *Cognition* **111**, 317–328. (doi:10.1016/j.cognition.2009.02.012)
 32. Smith K, Wonnacott E. 2010 Eliminating unpredictable variation through iterated learning. *Cognition* **116**, 444–449. (doi:10.1016/j.cognition.2010.06.004)
 33. Kalish ML, Griffiths TL, Lewandowsky S. 2007 Iterated learning: intergenerational knowledge transmission reveals inductive biases. *Psychon. B Rev.* **14**, 288–294. (doi:10.3758/BF03194066)
 34. Griffiths TL, Christian BR, Kalish ML. 2008 Using category structures to test iterated learning as a method for identifying inductive biases. *Cogn. Sci.* **32**, 68–107. (doi:10.1080/03640210701801974)
 35. Caldwell CA, Eve RM. 2014 Persistence of contrasting traditions in cultural evolution: unpredictable payoffs generate slower rates of cultural change. *PLoS ONE* **9**, e99708. (doi:10.1371/journal.pone.0099708)
 36. Laland KN. 2004 Social learning strategies. *Learn. Behav.* **32**, 4–14. (doi:10.3758/BF03196002)
 37. Aoki K, Wakano JY, Feldman MW. 2005 The emergence of social learning in a temporally changing environment: a theoretical model. *Curr. Anthropol.* **46**, 334–340. (doi:10.1086/428791)
 38. Aoki K, Feldman MW. 2014 Evolution of learning strategies in temporally and spatially variable environments: a review of theory. *Theor. Popul. Biol.* **91**, 3–19. (doi:10.1016/j.tpb.2013.10.004)
 39. Boyd R, Richerson PJ. 1988 An evolutionary model of social learning: the effects of spatial and temporal variation. In *Social learning: psychological and biological perspectives* (eds T Zentall, BG Galef), pp. 29–48. Hillsdale, NJ: Lawrence Erlbaum Associates.
 40. Feldman MW, Aoki K, Kumm J. 1996 Individual versus social learning: evolutionary analysis in a fluctuating environment. *Anthropol. Sci.* **104**, 209–213. (doi:10.1537/ase.104.209)
 41. Rogers A. 1988 Does biology constrain culture? *Am. Anthropol.* **90**, 819. (doi:10.1525/aa.1988.90.4.02a00030)
 42. Kandler A, Laland KN. 2013 Tradeoffs between the strength of conformity and number of conformists in variable environments. *J. Theor. Biol.* **332**, 191–202. (doi:10.1016/j.jtbi.2013.04.023)
 43. Henrich J, Boyd R. 1998 The evolution of conformist transmission and the emergence of between-group differences. *Evol. Hum. Behav.* **19**, 215–241. (doi:10.1016/S1090-5138(98)00018-X)
 44. Nakahashi W. 2007 The evolution of conformist transmission in social learning when the environment changes periodically. *Theor. Popul. Biol.* **72**, 52–66. (doi:10.1016/j.tpb.2007.03.003)
 45. Wakano J, Aoki K. 2007 Do social learning and conformist bias coevolve? Henrich and Boyd revisited. *Theor. Popul. Biol.* **72**, 504–512. (doi:10.1016/j.tpb.2007.04.003)
 46. Kendal JR, Giraldeau L-A, Laland KN. 2009 The evolution of social learning rules: payoff-biased and frequency-dependent biased transmission. *J. Theor. Biol.* **260**, 210–219. (doi:10.1016/j.jtbi.2009.05.029)
 47. Nakahashi W, Wakano JY, Henrich J. 2012 Adaptive social learning strategies in temporally and spatially varying environments: how temporal vs. spatial variation, number of cultural traits, and costs of learning influence the evolution of conformist-biased transmission, payoff-biased transmission, and individual learning. *Hum. Nat.* **234**, 386–418. (doi:10.1007/s12110-012-9151-y)
 48. Kempe M, Lycett S, Mesoudi A. 2012 An experimental test of the accumulated copying error model of cultural mutation for Acheulean handaxe size. *PLoS ONE* **7**, e48333. (doi:10.1371/journal.pone.0048333)